Measurement and Digitization of Phase Angle Modulated Sensor Data

by

Frank J. Leibly III

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Bachelor of Science

and

Master of Science in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1994

Author. . . . . .
Department of Electrical Engineering and Computer Science
Sept. 9, 1994

Certified by . . . . . .
Professor Donald E. Troxel
Professor
Thesis Supervisor

Certified by . . . . . .
Edmund Foster
Draper Staff
Thesis Supervisor

Accepted by . . . . . .
Frederic R. Morgenthaler
Chair, Department Committee on Graduate Students

Measurement and Digitization of Phase Angle Modulated Sensor Data

by

Frank J. Leibly III

Submitted to the Department of Electrical Engineering and Computer Science
on September 9, 1994, in partial fulfillment of the requirements for the degrees of
Bachelor of Science
and
Master of Science in Electrical Engineering

Abstract

A highly accurate digital architecture for demodulating low rate narrow band phase angle modulated sensor signals was developed. Phase resolution finer than half a clock LSB period of time has been achieved. The application of interest for such a readout is to demodulate and digitize signals from inertial instruments in a guidance system. Several demodulation architectures have been evaluated and discussed. A general treatment of phase digitization resolution and accuracy limitations has been presented, applicable to all digital phase measurement implementations that result in synchronization induced quantization errors, including the most commonly used digital and digitizing phase locked loops.

Thesis Supervisor: Donald E. Troxel
Title: Professor

Thesis Supervisor: Edmund Foster
Title: Draper Staff

Acknowledgments

I am grateful to the staff of Charles Stark Draper Lab for providing assistance with my studies through the majority of my undergraduate and graduate years at MIT. During a series of difficult years for the lab, the continuous support given to my research has been greatly appreciated.

In particular, I'd like to thank Ed Foster and Don Troxel for supervising my thesis; Ken Fertig, for hours of thought provoking and insightful discussions leading to an enlightening thesis topic; Sumner Brown and Bob England, who offered their assistance and supervised my work for several years at the lab; Ken Gerish and Dick Gardner for turning me loose into an extremely broad and fascinating topic area, and for showing faith in my judgment and abilities; all of the others at ETD and Draper as a whole who helped make my experience at the lab a pleasant and educational experience; and the folks at NFIC for understanding my need to set aside time from the company to complete my thesis work. Foremost, I am grateful to my family for their continued support.

This thesis was prepared at The Charles Stark Draper Laboratory, Inc., under United States Air Force Contract F 04704-88-C-0100. Publication of this thesis does not constitute approval by Draper or the sponsoring agency of the findings or conclusions continued herein. It is published for the exchange and stimulation of ideas.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

In inertial guidance systems, it is frequently necessary to measure and digitize the relative phase angle between certain signals. In particular, many types of accelerometers and gyroscopes generate signals which are phase modulated with respect to the guidance clock in response to accelerations and torques.

To acquire, track, and digitize the angle of these phase modulated signals in systems which have developed at The Charles Stark Draper Laboratory, a traditional method has been to use highly accurate analog phase locked loops. While these designs have proven to be extremely useful, an all digital phase readout and digitization scheme is sought for potential use in new systems. The reported research investigates the viability of digital phase readout and digitization techniques, and addresses some of the fundamental design decisions and limitations present in phase measurement and digitization of narrow band phase modulated signals.

Since a specific system is not being developed, the analysis herein is somewhat broad in scope, with an emphasis on developing a thorough understanding of digital phase measurement techniques and limitations that might play an important role in future instrument readout systems.

1.1 Guidance System Applications

In an inertial guidance system, instruments detect acceleration and torques, typically producing frequency and phase modulated signals. These signals must be demodulated, digitized, processed, and sent to a guidance computer, which integrates the information to determine an accurate estimation of location, velocity vector, and acceleration vector. Detailed gravity models are used to subtract out the gravity contributions to these measurements in the onboard guidance computer.

Several instruments which require phase angle measurement and digitization in this manner are the Specific Force Integrating Receiver (SFIR), the Pendulous Integrating Gyroscopic Accelerometer (PIGA), the Vibrating Beam Accelerometer (VBA), the Microwave Resonant Accelerometer (MRA) , the Interferometric Fiber Optic Gyroscope (IFOG), and several types of Laser Gyros. Each of these instruments produces an output frequency which is sensitive to linear or rotational acceleration along or around an axis, and is either used in or is under consideration for use in highly accurate guidance applications. In each case, the frequency/acceleration relationship is approximately linear, and signal phase can be used as a first order indication of velocity or angular rate along or around the instrument's input axis.

While the context of the reported research is inertial guidance, the results are applicable, in general, to any system where narrow band phase modulation is to be digitally recovered.

1.2 Phase Vs. Frequency Measurement

Since inertial guidance requires highly accurate velocity information, it is critical for instrument readout circuitry to accurately accumulate phase measurements to maintain a

good estimation of velocity. Since velocity can not be physically detected, instruments produce frequencies which are proportional to detected acceleration, or specific force. A technique for determining velocity from acceleration measurements is to periodically measure the frequency of an instrument, and integrate these measurements to get velocity data. A more natural and appropriate technique to make this measurement is to utilize the phase of an instrument's output signal as a continuous integral of signal frequency which is far more accurate than any scaled sum of discrete frequency measurements. Hence, measuring the phase of an instrument's signal is a more direct and accurate indication of velocity (or angular rate) then measuring its frequency. For this reason, the techniques discussed herein are evaluated for accurate phase angle measurement rather than frequency measurement, a subtle yet important distinction.

The importance of true phase measurements is particularly apparent when a perfect linear accelerometer is considered, where the frequency deviation of the device varies linearly with acceleration and is error free. In such a case, if the instantaneous phase of the instrument is known initially (along with an initial velocity), a phase measurement at any other time (past or future) results in a perfectly corresponding velocity measurement. This assumes, of course, that a full, unquantized, unwrapped phase word is available, i.e. the number of complete signal cycles which have occurred prior to the current unquantized partial cycle measurement are known.

1.3 Guiding Principles

Several guiding principles drive the consideration and analysis of digital phase measurement techniques for guidance applications. In particular, it is thought that when a specific digital phase readout approach is implemented, it could take on a simple and general enough form that it would offer applicability in a broad range of systems. Some

of the driving criteria behind the development of an general purpose phase modulated signal readout system are enumerated as follows.[1]

1) A single circuitry which can recover phase modulated signals over a very wide range of frequencies, while maintaining the performance of instrument specific readouts could facilitate future instrument and measurement system design efforts.

2) A modular system which shares common resource would be easier to design, deploy, and upgrade.

3) In particular, a semi-autonomous signal processing module component might be frequently upgraded and improved as integrated circuit signal processing power continues to rise, and costs continue to fall.

4) A primarily digital implementation would facilitate new applications and improvements with only software modifications.

5) A straightforward approach which is well modeled and understood will be robust and insensitive to circuit non-idealities.

6) A flexible readout system could be used in specific real time applications, in a test environment, or an adaptive system to optimally perform under various conditions.

A digital solution for accurate phase angle measurement is sought for several of the aforementioned reasons. In particular, a digital solution is thought to be a more

[1]Appendix D applies these principles for a digital phase readout to the potential development of a universal instrument readout system.

flexible component of a modular system, potentially allowing the same readout to be used for multiple instruments in the same system simultaneously, or in completely different systems. As digital processing power becomes increasing abundant and inexpensive, the ability for extensive real time computation and estimation without analog pre-processing becomes more feasible and attractive in instrumentation. Finally, digital circuitry is less sensitive in a radiation environment, making it preferable in certain applications of interest for such a readout.

1.4 The Digitizing Analog Phase Locked Loop

Several digital readout design implementations are considered for the demodulation and digitization of phase angle modulated sensor signals in a guidance system. An important component in existing systems is the digitizing analog phase locked loop shown in figure 1.1 below. This architecture has been successfully implemented for sensor readout in existing guidance systems, and serves as a good benchmark for comparing alternative readout implementations.[1]

[1]Fertig and Cox provide a detailed discussion of the digitizing analog phase locked loop, as a high speed I/O mechanism with real time computing capability. The strobing process is clearly depicted in several drawings included with their report.

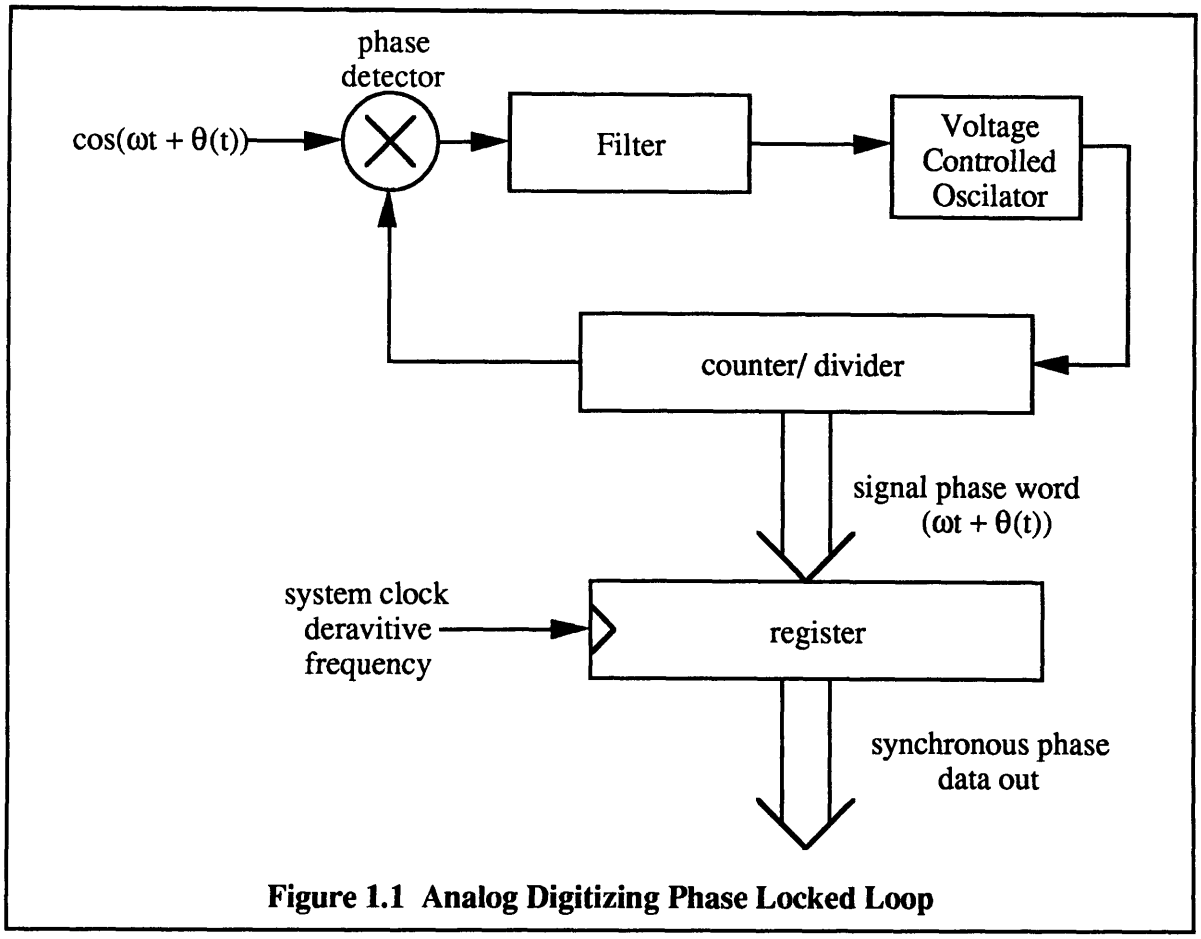

Figure 1.1 Analog Digitizing Phase Locked Loop

The operation of the analog digitizing phase locked loop is briefly explained as follows. The counted down high frequency VCO is synchronized by the feedback loop with the incoming signal, resulting in a binary representation of signal phase which can be strobed by the onboard clock to yield relative phase information. The quantization resolution of this approach is limited, in some fundamental way, by the VCO frequency, and corresponding number of bits in the feedback counter.

1.4 The Digital Clock Strobing Readout

An alternate approach suggested is a clock strobing architecture, as opposed to the signal strobing architecture described in section 1.3. In the Digital Clock Strobing Readout

(DCSR), the sensor signal zero crossings are used to strobe all the bits of the onboard clock[1], just as the onboard clock is used to strobe bits synchronized with the incoming signal in the digitizing phase locked loop configuration. The Digital Clock Strobing Readout is pictured in figure 1.2 below.

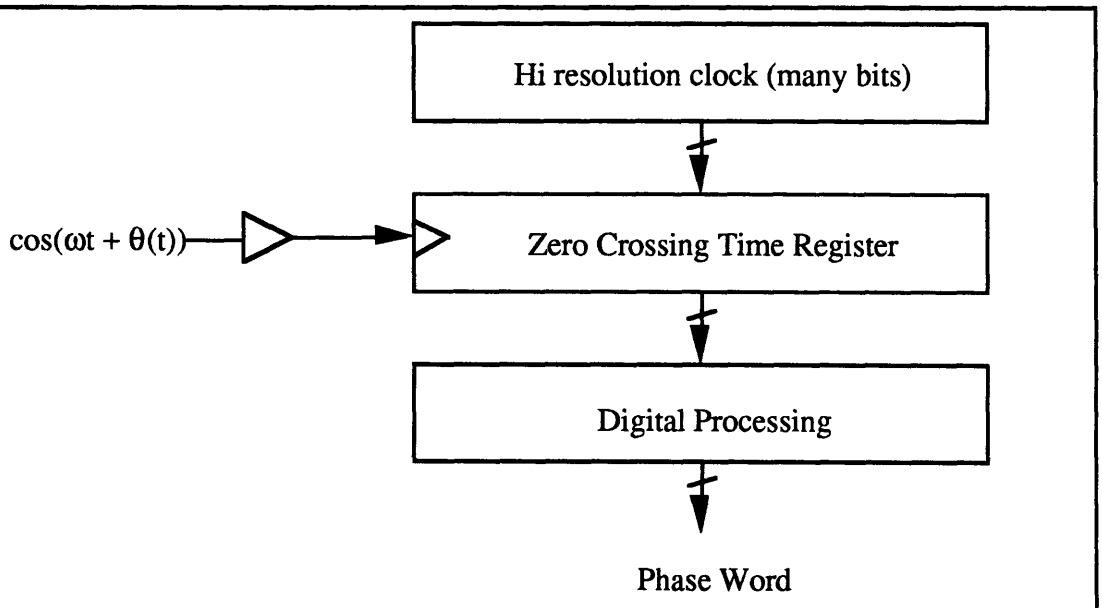

Figure 1.2 Digital Clock Strobing Readout

This DCSR generates "clock phase words" or "zero crossing times" each time the incident signal passes through zero voltage. The clock strobing process, analogous to feedback counter strobing for the analog digitizing phase locked loop, is pictured in figure 1.3 below.

[1]In a guidance system, not one, but many clock bits are available. Using techniques discussed by Comstock, the whole clock word maintains integrity, even over periods where the counter fails to operate for some number of cycles (disruption). This is accomplished by checking the whole time word against the phase of low frequencies generated by heterodyning techniques, and correcting the counter word when a difference occurs.

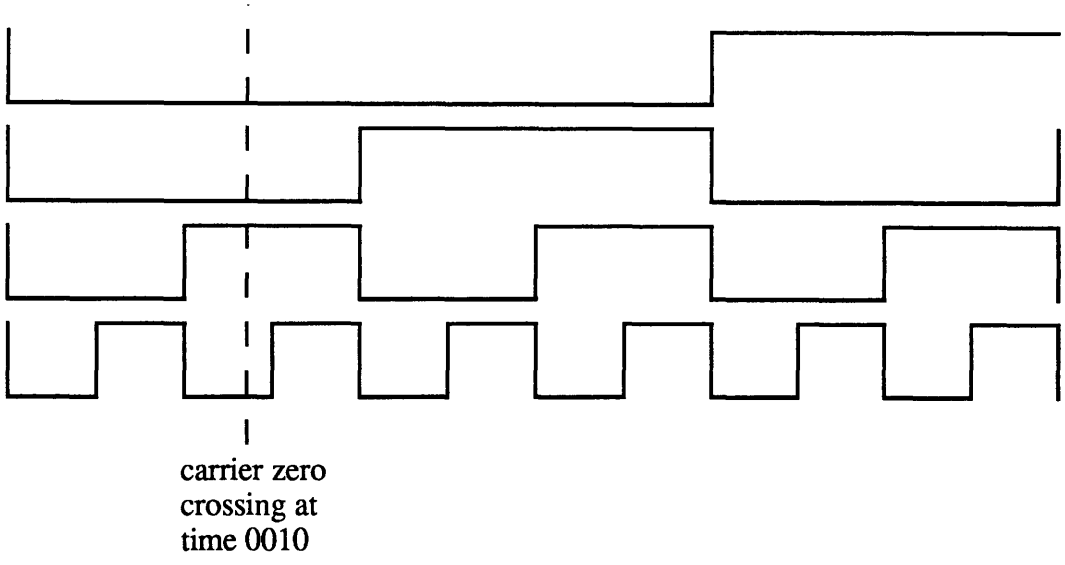

Figure 1.3 Clock Strobing Process

Several potential advantages are seen to the clock strobing approach. Consider, for example, that a high resolution clock already contains many bits. For this reason, it may not be necessary to generate higher frequencies as is necessary when a comparably low frequency signal phase is to be strobed. If in fact it is necessary to beat the clock up to higher frequencies, the steady frequency and phase of the clock might be tracked more reliably than a dynamically changing signal. A second related potential advantage is the ease with which this type of readout system could be used to demodulate several signals from different sources: the same high frequency clock may be strobed by an unlimited number of signals, resulting in an efficient use of resources.

1.5 The Delta Sigma Phase Locked Loop

Another architecture considered, in considerably less detail, is a mixed signal or completely digital phase locked loop, using Delta Sigma[1] modulation techniques to

[1]Delta Sigma Modulation is are also referred to in the literature as Sigma Delta Modulation.

ultimately obtain more bits of phase resolution than would otherwise be measured[1]. A potential implementation of the approach is indicated in figure 1.4 below.

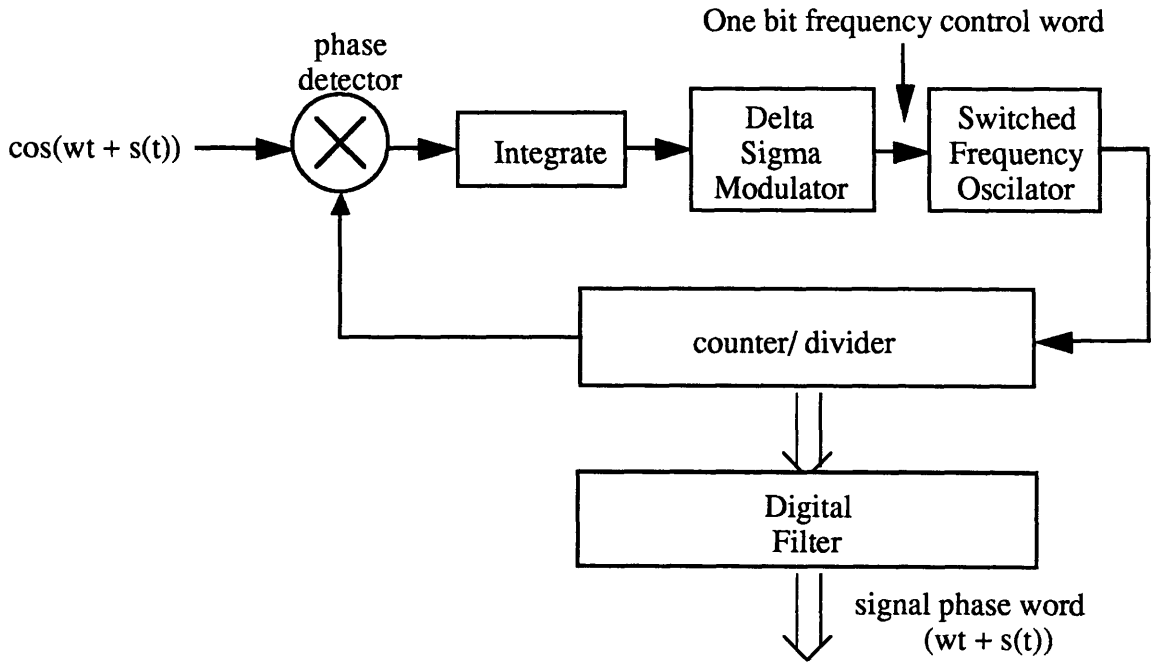

Figure 1.4 The Delta Sigma Phase Locked Loop

The rationale behind considering the Delta Sigma configuration is the possibility of shifting phase quantization noise due to finite oscillator or clock frequency beyond the signal spectrum, so that it can be subsequently removed. In baseband A to D and D to A converter implementations, Delta Sigma configurations are implemented for exactly this purpose.

1.6 Summary

Chapters 2 through 7 examine the process of phase readout and digitization of narrow

[1]The Delta Sigma Phase Locked Loop is considered for the case of a digital phase detector in Chapter 6, and with an analog phase detector in Appendix C ,which discusses hybrid analog/digital phase locked loop techniques.

band phase modulated signals, in the context of inertial instrument readout. Particular emphasis is placed on the motivation, benefits, and limitations of the Digital Clock Strobing Readout, which is shown to offer the same or better resolution as an Analog Digitizing Phase Locked Loop run at comparable frequencies. In addition to examining particular phase readout implementations, an attempt is made to generalize the process of phase readout and digitization, and to expose the fundamental issues and limitations which occur across a broad spectrum of implementations. A fundamental question considered is the ability to achieve phase resolution beyond that of the highest available frequency or bit which is strobed or used to synchronize phase error detection. In any case, this is the phase resolution limit imposed for a single phase or phase error measurement. The notion of averaging successive phase measurements to obtain a finer representation of phase is evaluated in detail.

Chapter 2, Phase Readout Functionality, breaks down the process of phase readout and digitization into a set of semi-independent tasks that are performed by any phase readout and digitization scheme. The goal of Chapter 2 is to develop a better understanding and a working vocabulary for analyzing and evaluating the specific phase readout and digitization approaches proposed and considered in the thesis.

Chapter 3, Phase Locked Loops, discusses the functionality, performance, and limitations of analog and digital phase locked loop implementations for phase readout and digitization. In particular, it is noted that both digital and analog digitizing phase locked loops suffer from synchronization induced quantization errors which limit their performance. It is also noted that the implementation of finite order loop filters induces dynamic tracking errors in both types of implementations.

Chapter 4, The Digital Clock Strobing Readout, explains and discusses the DCSR architecture and evaluates it's performance and limitations. Specifically, the difference and similarity between clock strobing and signal strobing are explained. The

fundamental measurement of interest is shown to be the *phase relationship* between two signals, rather than the absolute phase of either. As a means of measuring this phase relationship, the advantages and disadvantages of clock strobing are evaluated. In particular, the resolution limitations of the readout architecture are determined in the static case, and these results are used to set a bound on dynamic resolution for narrow band phase modulated information. The ultimate phase resolution which may be obtained is related to the signal frequency deviation ratio, the phase noise content, the signal bandwidth, and the ratio of clock to carrier frequencies. It is shown that in some cases, minimum resolution far beyond the clock's resolution may be achieved by averaging phase measurements, while in other cases resolution is limited to a fractional clock cycle determined by the clock to carrier frequency ratio. Several algorithms for the processing of DCSR measurements to yield filtered and demodulated phase words of desired resolution are proposed.

Chapter 5, Simulation of the Digital Clock Strobing Readout, describes several sets of computer simulations which have been performed to test and further characterize the fundamental phase resolution limitations of the DCSR. The minimum resolution limitations described in Chapter 4 are confirmed, and several additional phase measurement resolution metrics are considered. Simulations of the readout resolution in response to dynamic phase modulated signals demonstrates how the static resolution characteristics of the readout may be used to predict dynamic resolution for bandlimited signals.

Chapter 6, DCSR Analysis Applied to Other Phase Readout Approaches, describes how the analysis and simulation presented for the DCSR in Chapters 5 and 6 apply, in general, to the class of phase readout and digitization techniques which result in synchronization induced quantization errors. This class of readout systems includes the broadest family of digital phase locked loops, analog digitizing loops, and the Digital

Clock Strobing Readout. The limitations of the Delta Sigma Phase Locked Loop with a digital phase detector are included in this class as well. Using the analysis from Chapter 4, it is shown how the resolution of the analog digitizing phase locked loop may be improved in certain cases beyond the VCO resolution by appropriately modifying the feedback counter length and strobing rate, in conjunction with averaging the resulting signal phase measurements.

Chapter 7, Concluding Remarks, offers an overview of the notable conclusions from the thesis. Appendix A details the spreadsheet simulation techniques used to perform the experiments reported in Chapter 5. Appendix B discusses a number of hybrid analog/digital clock strobing techniques which may offer considerably improved performance beyond the resolution limitations of the DCSR. Appendix C discusses hybrid analog/digital phase locked loop techniques which overcome the limitation imposed by synchronization induced quantization errors. By using an analog phase detector, these readouts, including the Delta Sigma Phase Locked Loop, obtain a digitized phase word with resolution significantly higher than the clock's resolution. Appendix D discusses a Universal Instrument Readout, implemented using the DCSR in conjunction with a highly accurate voltage measurement system such as a voltage to frequency converter.

Chapter 2: Phase Readout Functionality

Before considering specific approaches for phase readout and digitization in detail, it is interesting to note that each approach must share certain functionalities in common, and it is the implementation of each of these functions which distinguishes each design alternative. This chapter breaks down the problem of digitizing phase modulated signals into a set of specific functions which are performed in each case, and takes a brief look at some of the ways each function might be realized. The basic functions of phase readout and digitization considered in this chapter are:

- synchronization of incoming data to an onboard clock
- demodulation (carrier removal)
- phase noise filtering
- interpolation/extrapolation
- memory (if signal or circuitry cuts out for some amount of time)
- data pre-processing
- quantization of phase measurement
- filtering of quanitation noise
- data rate conversion

It will be interesting to note, when considering the individual phase readout designs in later chapters, how the functions need not be implemented in a particular order, nor is it necessary for the tasks to be done independently.

2.1 Data Synchronization

In the general case, a phase modulated signal is not synchronous with the digital circuitry which will process or ultimately receive the digitized measurements. If in fact the incoming signal is a phase modulated sinusoid, than it can be considered continuous data which must be sampled (synchronously or asynchronously). If the sinusoid is squared up, it becomes a 1 bit asynchronous digital signal. Either way, at some point in the phase readout and conversion process, it is necessary to provide signal information synchronous with the clock driving the digital circuitry used for processing or storing the data. Hence a general function performed in any phase measurement and digitization system is data synchronization, as depicted in figure 2.1 below.

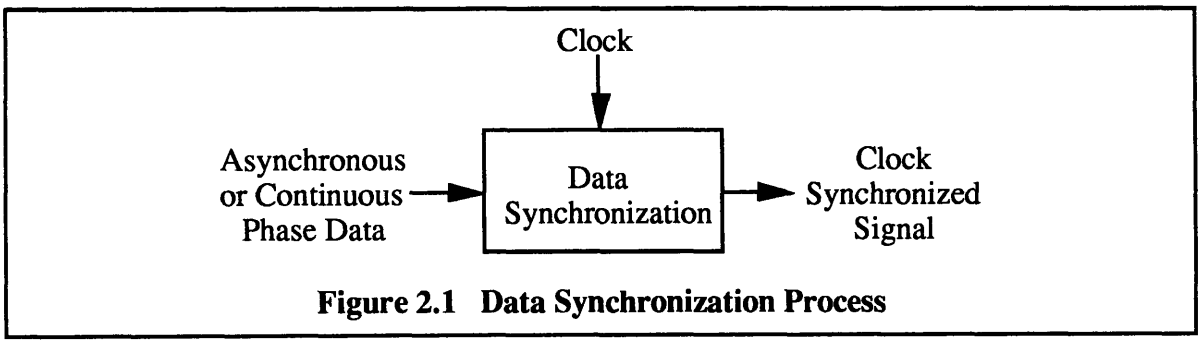

Figure 2.1 Data Synchronization Process

A simple, and straightforward way to perform this synchronization is to sample the incoming phase modulated signal synchronously periodically. These samples can be digitized and processed to ultimately retrieve the desired measurements. This approach is depicted in figure 2.2 below.

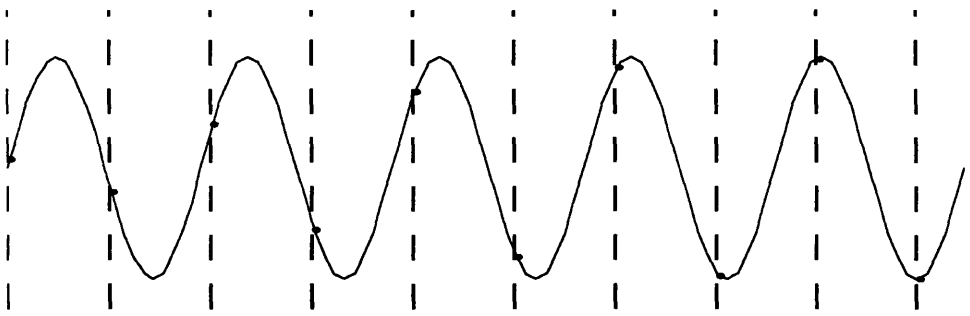

Figure 2.2 Synchronization by Periodic Sampling

While this process is typical for synchronization with baseband analog information, it makes less sense when dealing with phase modulated signals. There is a class of digital phase locked loops which performs synchronization in this exact manner, referred to by Lindsey and Chie[1] as Nyquist Rate Digital Phase Locked Loops, but it is not a common approach.

A more common way to synchronize phase modulated information is to synchronize a one bit squared signal by latching the data with a high speed flip flop. This method of synchronization is depicted in figure 2.3 below.

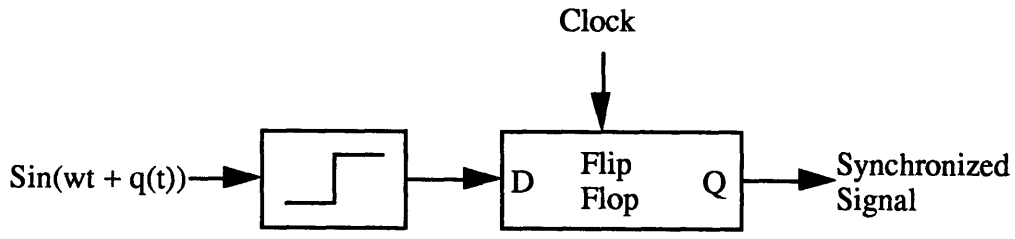

Figure 2.3 Synchronization by Latching

The obvious result of latching is an induced phase change in the resulting signal. The resulting synchronized signal no longer has the same phase characteristics as the original

[1]In "A survey of Digital Phase-Locked Loops," Lindsey and Chie identify four different classes of digital phase locked loops, based on the type of phase error detection used.

asynchronous signal. This synchronization induced phase offset, depicted in figure 2.4 below, is a dominant source of measurement error for phase readout implementations which use this method of synchronization, and those that behave the same way using different components.

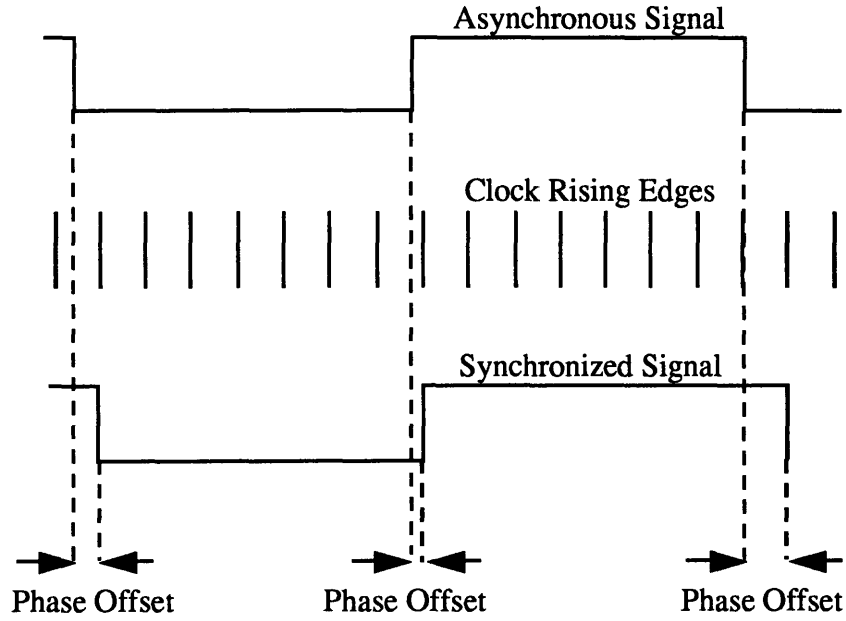

Figure 2.4 Synchronization Induced Phase Quantization Errors

While the implementation in figure 2.4 generates data which is synchronous with the onboard clock, the samples are still not uniformly spaced in time, requiring additional processing to attain periodic synchronous data.

Synchronization plays a key role in determining the ultimate phase resolution which can be achieved for a particular readout implementation. The resolution limitations discussed in Chapters 4, 5 and 6 will be attributed specifically to synchronization induced quantization errors..

2.2 Demodulation

Demodulation, or extracting the phase of an incoming signal relative to some nominal signal's phase, is the most fundamental burden of the circuits considered. As with synchronization, demodulation can be effectively accomplished in a number of ways.

Given periodic samples of a phase modulated sinusoid, discussed as an option in the section 2.1, an arcsin computation could effectively recover the original signal instantaneous phase at each of the sampling times. The word "demodulation," however, implies that some nominal frequency component needs to be subtracted off of the instantaneous phase, in order to yield the desired relative phase information. This direct approach is depicted in figure 2.5 below.

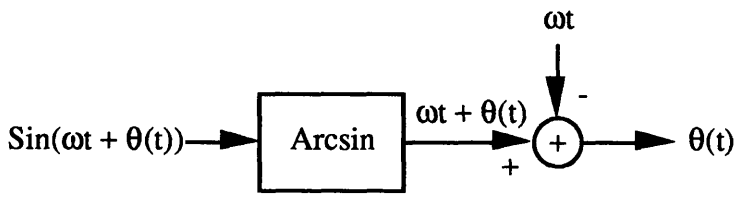

Figure 2.5 Direct Demodulation Approach

A more traditional demodulation approach is to multiply the incoming signal, or a digital representation of it, by a quadrature version of the nominal frequency and phase reference carrier, yielding a average product equal to the sin of the phase difference. In this case, the multiplication and lowpass filtering effectively subtract off the nominal frequency term, this time prior to an arcsin computation.

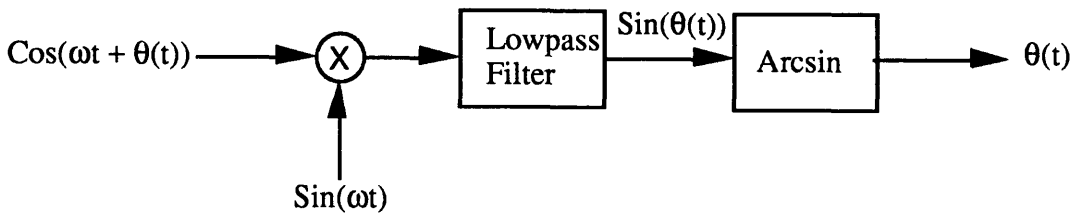

Figure 2.6 Demodulation by Quadrature Multiplication

If the phase difference in the above demodulation scheme is very small (where the radian slope of the sin function is approximately 1), then the arcsin function can be eliminated. This is a good approximation in a closed loop system such as a phase locked loop[1].

Regardless of the specific demodulation implementation, any demodulation method needs to consider both the extraction of phase data, and the conversion of phase to relative phase signals.

2.3 Phase Noise Filtering

Because signals in the real world are corrupted with noise, it is generally important that efforts are made to reduce the sensitivity of measurements to noise. Hence, some sort of filtering is typically done to remove the effects of signal noise from phase measurements.

The noise sensitivity of phase measurements depends partially on the point in time which the phase is sampled. It can be easily demonstrated that the phase of a sinusoidal phase modulated signal is most sensitive to noise at times when the absolute phase angle is near 90 or 270 degrees. At these phase values, signal amplitude changes very slightly with respect to phase; i.e. a small amount of additive noise can result in a large measured phase error. Likewise, phase measurements are least sensitive to noise at

[1]In this picture, the assumption is made that the lowpass filter removes all carrier frequency harmonics. This approximation is useful for understanding the intended principle of operation of this demodulation method, but the effects of non-ideal filtering must be considered when evaluating the performance of phase locked loops which demodulate in this manner to produce error signals.

0 and 180 degrees, when the signal value is changing maximally with respect to phase. Formally stated, the sensitivity of phase measurement to additive noise is inversely proportional to signal derivative with respect to phase, having a value of infinity at 90 and 270 degrees, and minimum values of 1 at 0 and 180 degrees.

One of the most significant improvements in the removal of phase noise is accomplished merely by amplifying and limiting the instrument output before much noise corrupts the signal. The almost squared up signal (with finite rise and fall times) is extremely immune to noise since the slope of the signal at zero crossings is very large.

Phase noise filtering can be done independently, or in conjunction with other filtering, i.e. of quantization noise. Or quite possibly, the phase noise could be filtered at multiple stages within the circuit.

2.4 Interpolation/Extrapolation of Phase

Interpolation or extrapolation of phase is necessary in a phase readout whenever phase measurements are desired at times other than when samples are actually taken. With a digital phase readout where samples effectively occur at signal zero crossing times, some sort of interpolation or extrapolation of phase must take place in order to generate phase information at times other than the signal zero crossing times.

A common approach to extrapolation is to use a phase locked loop. Extrapolating phase between zero-crossing times is a primary motivation for using a phase locked loop as a component of a phase readout system. The loop accomplishes phase extrapolation by dynamically tracking the phase and phase rate of the signal, and using this rate to extrapolate phase at other times, as indicated in figure 2.7 below.

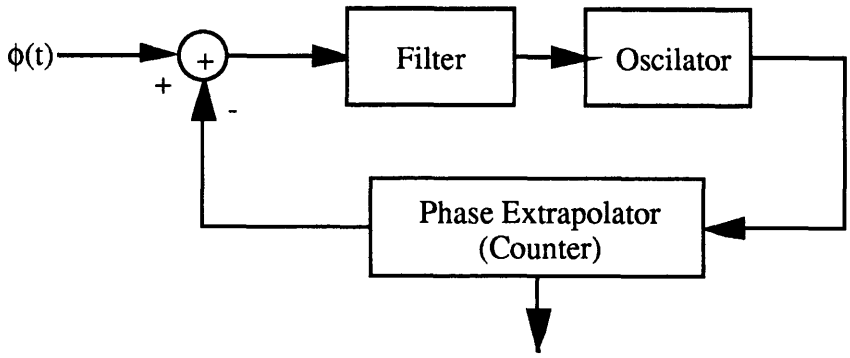

Figure 2.7 Phase Locked Loop Extrapolation

In addition to using different extrapolation techniques, phase interpolation can be used if phase measurements are not required in real time. Given a sequence of phase measurements, interpolating methods to generate values at intermediate points depend on what type of signal and noise model are used.

Consider the case where phase measurements are equally spaced and it is known that phase variations are restricted to a lowpass bandwidth less than half the sampling rate. In this case, interpolation can be effectively accomplished by padding the signal with an arbitrary number of zero's, followed by a low pass filter and gain stage. A factor of two fixed rate interpolator is diagrammed in figure 2.8 below.

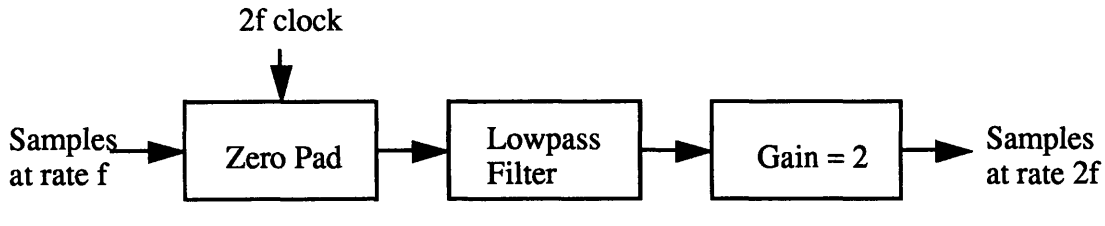

Figure 2.8 Interpolation by Upsampling

The interpolation/extrapolation question is in fact complicated by the fact that phase data generated from signal zero crossing times, once again in the case of the

squared carrier, does not arrive at a specific fraction of the onboard clock rate, nor does it arrive at fixed intervals whatsoever for a changing signal value. Hence under these circumstances, extrapolation or interpolation methods must account for variable rate data, as depicted in figure 2.9 below.

Figure 2.9 Complex Interpolation

2.5 Memory

Phase readout memory is particularly important in guidance applications, where integrated information is the desired quantity. If a signal disappears for some amount of time, the circuitry should be able to determine what happened and accurately compensate it's measurements. Hence the readout must *remember* previous phase measurements and use them to determine whether or not it believes current measurements. If not, it must decide what measurement to report and integrate, and at what point to correct integrated quantities, as pictured below.

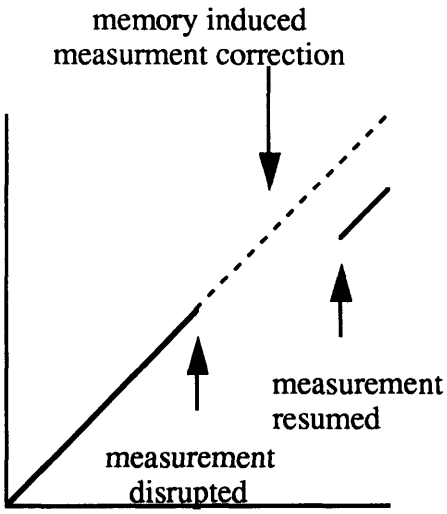

Figure 2.10 Disrupted Signal and Phase Readout Compensation

The actual implementation of memory in a readout system is quite specific to the particular readout implementation and is addressed individually with each of the designs considered. In each case, however, readout circuitry must somehow be able to accurately integrate measurements through times when the phase signal is disrupted, and adjust for measurement offsets resulting from such periods.[1]

2.6 Data Pre-Processing

Each of the other sections in Chapter 2 addresses fundamental issues of phase measurement circuitry. Data pre-processing is not essential, but is an important function of many phase readout designs and is considered for this reason.

Quite possibly, phase measurement data is taken at a relatively high rate compared to the rate which it is actually required. This high rate makes certain processing possible which could not be accomplished given data at the lower rate. Inclusion of this

[1]Comstock's techniques for recovering the state of an interrupted digital counter relate directly to memory implementation in both phase locked loop and clock strobing implementations.

processing into the measurement and readout system is a logical option. This processing can potentially be heavily integrated into the other aspects of the measurement system and accomplished with a minimal degree of added complexity.

An example which illustrates the potential to integrate signal preprocessing into a phase measurement system is the removal of periodic errors using a phase locked loop. Errors which occur periodically as a function of phase angle can be easily removed at a high rate with a lookup table at the output of the PLL, as shown in figure 2.11 below.

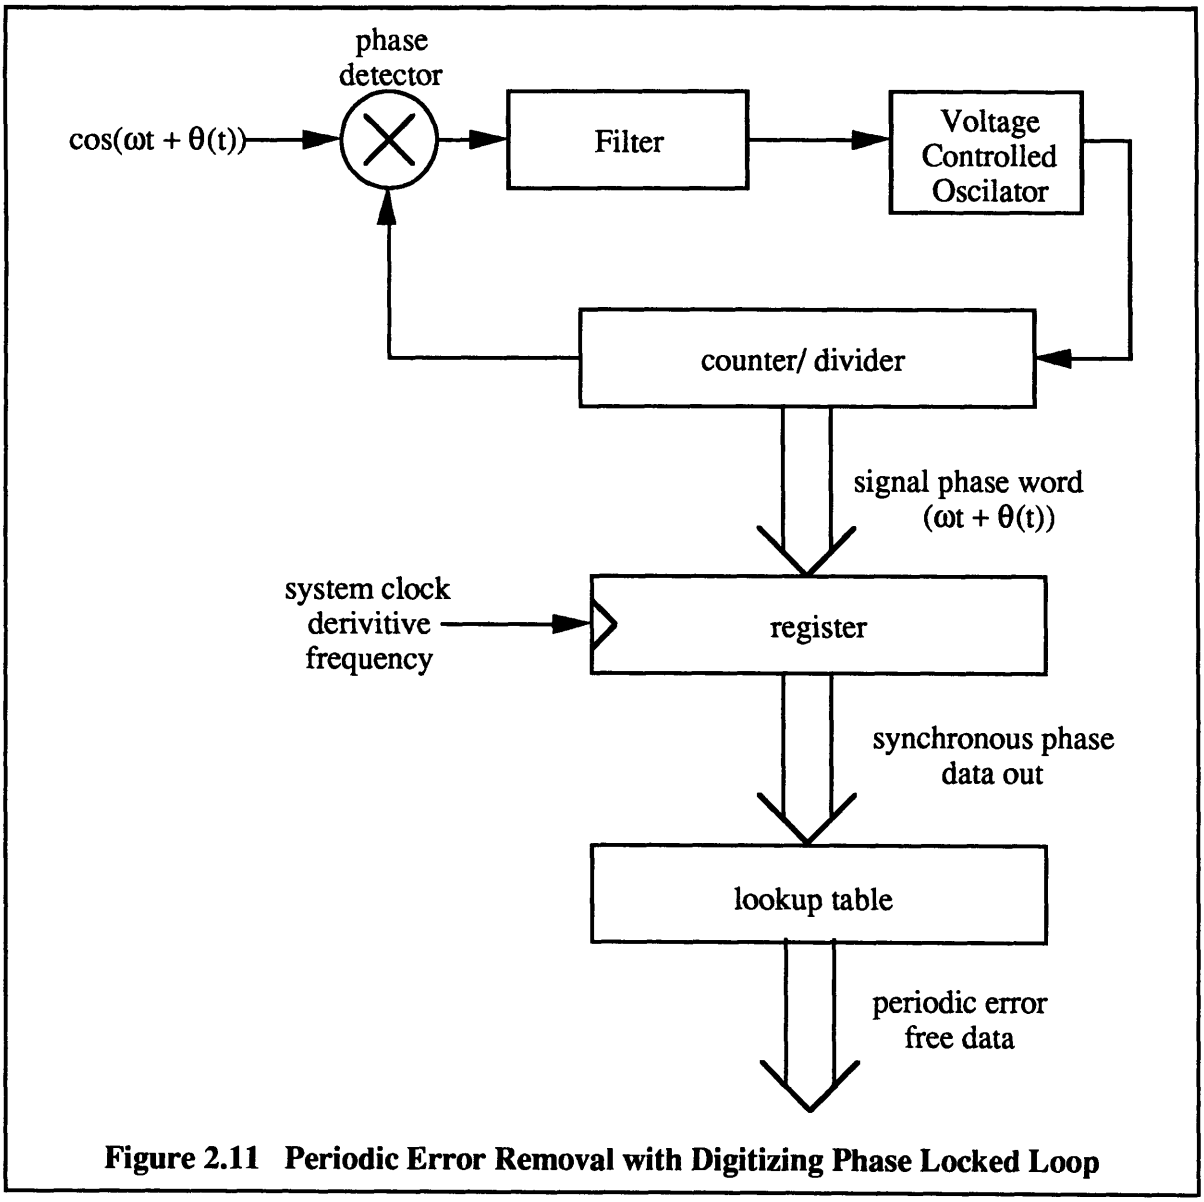

Figure 2.11 Periodic Error Removal with Digitizing Phase Locked Loop

Removing this error at later stages may be substantially more complicated and less reliable depending on what is done with the data subsequent to phase digitization.

2.7 Quantization Of Phase Measurement

Given the digital nature of computers, phase measurements must be represented with discrete numbers as well as in discrete time. For this reason, each phase measurement must be quantized to some digital word with known resolution.

Typically, quantized measurements destined for a computer originate from an analog voltage or current. In this case, quantization typically follows a greater than or equal to rule, whereby a continuously variable parameter (i.e. voltage) maps to a finite number of quantized numbers. Each quantized number represents a signal which was in reality, greater than or equal to some threshold and less than the next threshold. Typically these quantization intervals are derived from evenly distributed threshold voltages within an appropriate range for a given signal, i.e. 0-10v. This simple quantization scheme is pictured in figure 2.12 below.

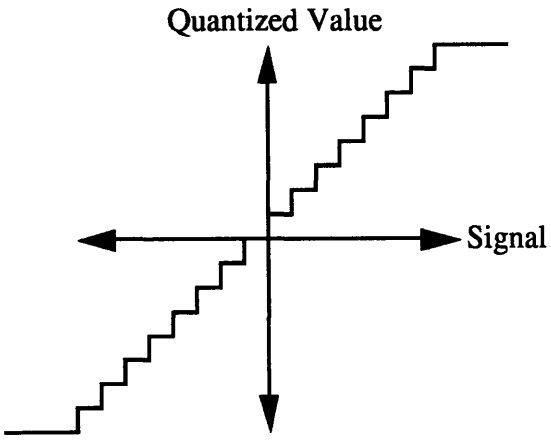

Figure 2.12 Fixed Level Quantization

In the case of phase measurement, quantization is complicated by the fact that establishing a set of reference phase quantization intervals requires an infinite number of quantization levels since there is no limit of the range of phase deviation. Also, there is no zero reference phase for a carrier which is not synchronous to the onboard clock. In such a case, quantization may establish the bounds of a particular phase measurement but the bounds may change each time a new measurement is taken.

The result of the changing quantization range is that not all measurements are the same. In fact, two measurements of the same signal phase could produce different quantized values. This affects the analysis of quantization noise as described below

2.8 Filtering Of Quantization Noise

The quantization resolution of an analog measurement is not necessarily limited by the number of available bits to represent the measurement, but by the accuracy of the converter. In some cases it may be possible to "improve" quantized measurements by "unquantizing" and "requantizing" with more bits of resolution. This can be accomplished by digitally filtering the quantized data, with zero's initially appended as least significant bits. The filter may be optimized for a particular signal and quantization noise model.

In the case depicted above, statistical properties of quantization noise can be derived given fixed quantization intervals. With variable quantization intervals, the results of this type of analysis need to be reformulated to optimize the performance of the filter given the nature of the quantization.

2.9 Data Rate Conversion

A final task of phase readout and conversion circuitry is to convert measurements to a rate demanded by external circuitry. This would in general be a fixed rate. Clearly, data rate conversion is closely tied to the synchronization and interpolation issues previously addressed.

A typical data rate conversion task might be to down sample fixed rate measurements subsequent to digital processing and send them at a slow fixed rate to the guidance computer. In this case, high frequency components of the data must be pre-filtered to avoid aliasing, and then the data can be converted to a lower rate. The simple and common case of integer multiple rate conversion is accomplished merely by throwing out most of the samples (after pre-filtering) as diagrammed in figure 2.13 below.

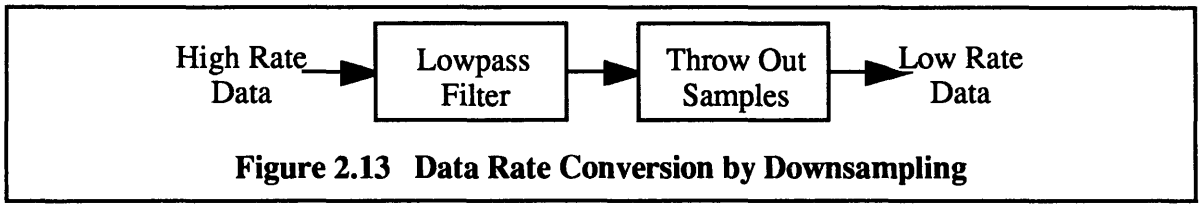

Figure 2.13 Data Rate Conversion by Downsampling

A more complex problem may be to convert the data to a rate which is not an integer fraction of the original rate. In this case, one approach may be to first up sample the data to an appropriate integer rate, and then down sample rate described above. The added stage of filtering provides the necessary interpolation to effectively report data at times when it was originally not available, though at a lower rate.

2.10 Conclusions

The phase measurement and digitization process generally involves a number of semi-independent processes, as described in each of the previous sections. While each process need not be represented by a specific block in the phase readout system, the ultimate design of the readout is determined by a set of decisions regarding how each of these functions is to be implemented, which functions will be lumped together, and in what order the operations are to be performed.

Chapter 3: Phase Locked Loops

A common approach to instrument phase readout and digitization utilizes a phase locked loop to implement many of the functions summarized in Chapter 2. Many of the systems which have been developed and used at Draper Laboratory have relied heavily on the use of phase locked loops for phase readout and digitization.

3.1 The Analog Digitizing Phase Locked Loop

An important and well understood phase readout architecture is the analog phase locked loop, which is easily implemented in a fashion which results in phase digitization. A simplified diagram of the analog phase locked loop is shown again in figure 3.1 below.

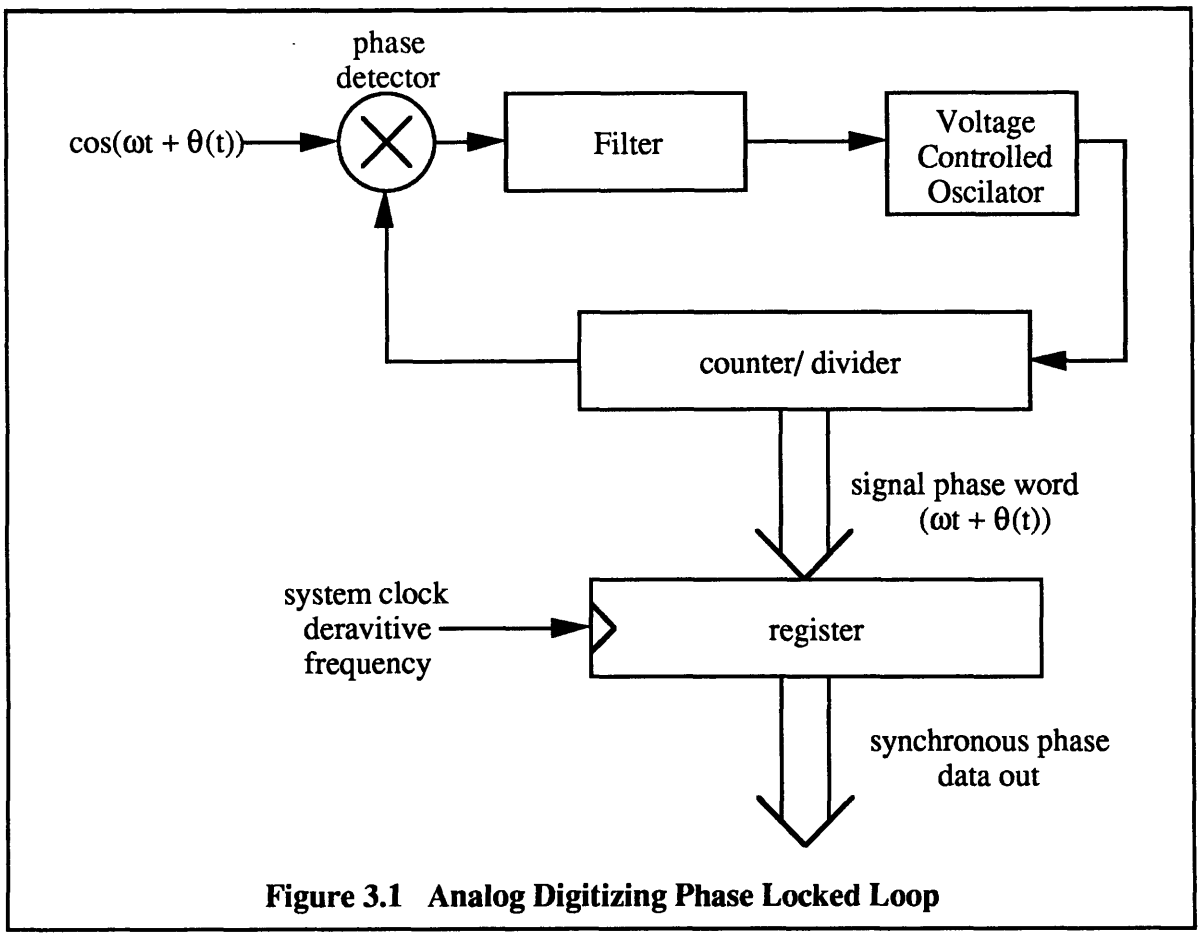

Figure 3.1 Analog Digitizing Phase Locked Loop

The above system locks a VCO frequency to some power of two multiple of the input frequency. The output of the counter divider is now a binary digitized version of total signal phase (wt + s(t)), which can be processed to extract the desired phase signal.

3.1.1 Functionality of Analog Digitizing Phase Locked Loop

The analog phase locked loop architecture accomplishes several functions simultaneously, making it a desirable component in many phase readout and digitization implementations. The functionality implemented using the analog digitizing phase locked loop combined with subsequent processing is depicted in figure 3.2 below.

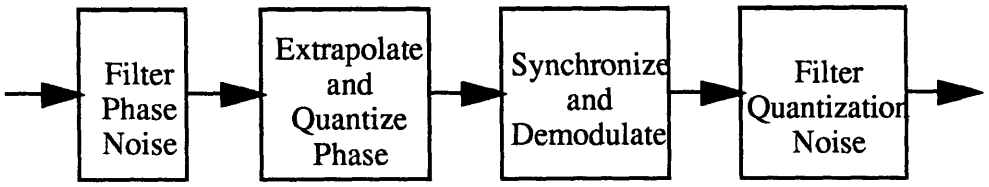

Figure 3.2 Functionality of Analog Digitizing Phase Locked Loop

Phase noise is removed by the loop filter, which is a lowpass filter designed to track the gradually changing phase of the signal, and eliminate higher frequency phase noise. With an integral term, the loop filter also provides memory for the system. Should the signal disappear at the input of the loop, the VCO will hold it's frequency (phase rate) reliably for some amount of time before a substantial phase error results. When the signal is recovered, even though many cycles may have been missed, the loop will have effectively tracked the missing cycles and resume tracking as if the signal had never been lost.

Phase extrapolation and quantization is accomplished by the VCO and feedback counter combination. Ideally, the counter performs ideal bands interpolation of phase between carrier zero crossings, though in practice it is actually an extrapolation. The number of phase quantization levels is determined by the counter length used to generate the feedback signal, which in turn is directly related to the VCO center frequency. If a single counter bit is used, four quantization levels are achieved, represented by the two bits comprising the VCO output and the counter output (at half the frequency of the VCO). For an n bit counter, n+1 bits of quantization resolution are attainable with a single VCO.

Synchronization is accomplished by strobing the feedback counter, which is presumably synchronized with the VCO and signal. Both synchronization and

demodulation[1] (carrier removal) can be accomplished simultaneously by strobing the feedback counter at appropriate times. If the feedback counter is strobed at a different rate than the carrier nominal frequency, then the carrier phase component must be subtracted off in a second stage following the counter strobing.

The phase locked loop architecture, in general, lends itself nicely to certain preprocessing tasks[2]. In situations where computer processing time is limited, handling pre-processing tasks directly in the phase locked loop can efficiently, and at a high rate, remove certain errors and compensate for certain conditions while freeing valuable compute time for other tasks [3].

The prospect of filtering out quantization noise from the phase locked loop's digitized phase output exists, but may in practice be limited to the extent that the noise bandwidth falls within the signal bandwidth itself. Since the quantization noise bandwidth is not arbitrary, but in fact tied inextricably to the rate and phase of the signal which strobes the counter, it may be the case that strobing to demodulate and strobing to shape quantization noise are conflicting tasks. For example, if the unmodulated carrier signal itself is used to strobe (and hence demodulate) the counter phase output, and the modulated phase error itself is fixed with respect to time for some interval, during this interval the phase quantization error will reside at DC in the spectrum, and be inseparable from the measurement itself.

Data rate conversion is an implicit function of the phase locked loop, which can provide data at any rate the feedback counter is strobed. The high rate at which data is available from the phase locked loop is attractive for additional data preprocessing

---

[1]While demodulation also refers to the process of multiplying the fedback signal times the carrier to create an error, I speak of demodulation in the sense that the carrier component of the signal is clearly present in the counter's phase measurement, and must be removed to recover the phase term alone.

[2]Fertig and Cox refer to a phase locked loop as having "real time computing capability" and offer several examples of how phase locked loops can be used to remove instrument errors and perform certain signal processing tasks.

[3]Fall's thesis discusses a number of resolver based instrument errors , the effect of which may be reduced with considerable digital processing of phase measurements.

beyond what is performed in the loop itself.

3.1.2 Loop Tracking Error

While the ideal phase locked loop would track bandlimited phase fluctuations without error, realizable loops do not perform ideally. In practice, the loop filter for an analog phase locked loop has a finite order, n, with a typical loop order n=2. Because of the finite filter order, sinusoidal phase modulations rare tracked with a certain amount of phase shift. The resulting loop error, plotted vs. phase modulation frequency, has a slope of 20n dB per decade upwards in the frequency plane toward an asymptote of 1 beginning where the modulation frequency equals the carrier frequency, as depicted in figure 3.3 below.

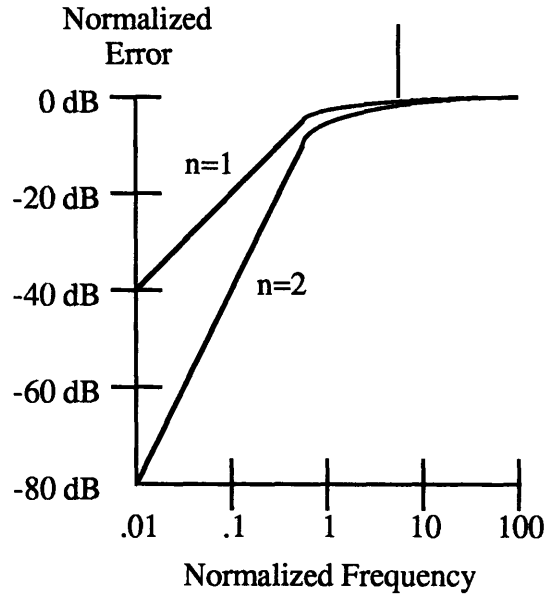

Figure 3.3 Loop Dynamic Tracking Error

Hence, for a second order loop with a modulation frequency equal to 1/100 times the carrier frequency, a resulting tracking error of 1 part in 10,000 (cycles) can be expected.

3.1.3 Limitations of The Analog Digitizing Phase Locked Loop

The primary limitation of the analog loop is flexibility. This is a result of the analog components, which are not programmable. The VCO limits the loop to operate within a narrow range of frequencies. Hence an analog phase locked loop designed and used for a particular signal will not be suitable for other applications without hardware modifications. In addition, design of an APLL is well documented, but nontrivial when due to stability and acquisition issues.

To track multiple phase modulated signals in a system, each signal requires it's own phase locked loop. Hence, with the desire to sharing resource, an analog phase locked loop offers no assistance. At best, digitized phase signals may share signal processing hardware.

The resolution limitations of the analog loop relate directly to the number of quantization levels defined by the VCO and feedback counter[1]. While zero DC loop error may be achieved when no noise is present, digitized phase is quantized as previously discussed, and not in such a way that filtering will necessarily remove quantization noise. This quantization error is in addition to any loop tracking error which exists, as discussed in section 3.1.2. A more complete discussion of the ultimate resolution which can be attained from the analog digitizing phase locked loop follows from the Digital Clock Strobing Readout analysis in Chapter 4

[1]Several ways to improve phase quantization resolution without using higher VCO frequencies but instead using additional VCO's are discussed in the Hypha Report. These approaches use multiple loops, increasing hardware and reducing flexibility in return for improved quantization resolution, and are based implicitly on the techniques demonstrated in section 6.2

3.2 Digital Phase Locked Loops

A similar approach to phase readout and digitization, modeled after the analog phase locked loop, but realizable with only digital components is the digital phase locked loop[1]. A typical digital phase locked is shown in figure 3.4 below.

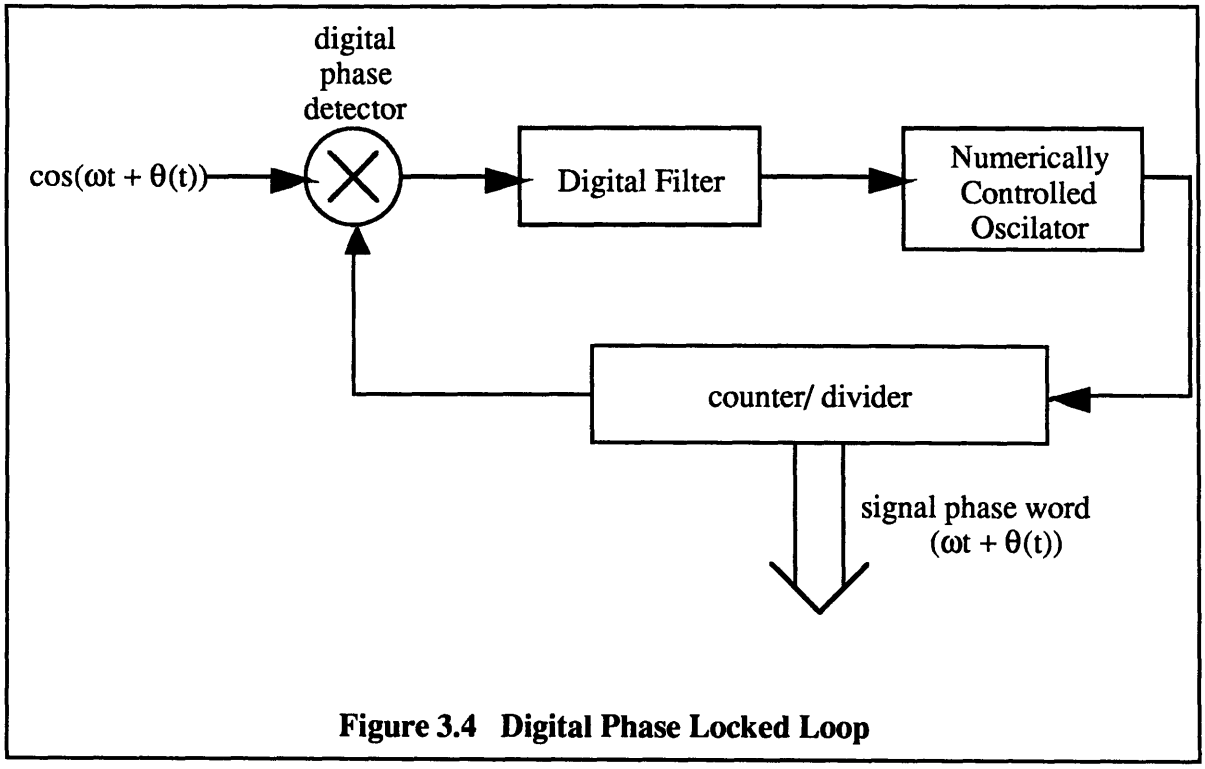

Figure 3.4 Digital Phase Locked Loop

3.2.1 Functionality of Digital Phase Locked Loops

The operation and functionality of the digital phase locked loop is very similar to the analog loop described above. The key difference is that all the components are now digital. This has several profound affects on the loop's performance.

---

[1]Bell, Cox, and Hanks discuss the design and use of digital phase locked loops for resolver processing. Cox (Draper Laboratory) received several patents for his digital phase locked loop design work. Other discussions of the use of digital phase locked loops for phase angle demodulation are found in the papers of Gill and Gupta, Lindsey and Chie, Natali, Osborne and Weinberg.

To begin with, the digital circuit can only cope with digital signals, hence the input signal, $\cos(\omega t + \phi(t))$ is typically squared up with a comparator. Another approach is to A to D convert the input signal. The latter approach is necessary when phase measurements are required at greater than twice the carrier frequency. This is not a typical situation, and is certainly not required for demodulation and digitization of narrowband phase modulated signals.

With a digital carrier (squared), a phase error is measured for each zero crossing of the 1 bit carrier by measuring the time between the zero crossings of the carrier and fed-back signals. This time, typically measured with a counter, is digitally filtered and used to control a numerical oscillator, typically another counter. The loop is designed such that when no phase error is present, the "numerical oscillator" counter keeps counting at a constant rate producing a rollover bit at the carrier frequency. This carry over bit is fedback to the carrier for demodulation (error measurement). The functionality of the digital phase locked loop, in terms of the tasks defined in Chapter 2, is depicted in figure 3.5 below.

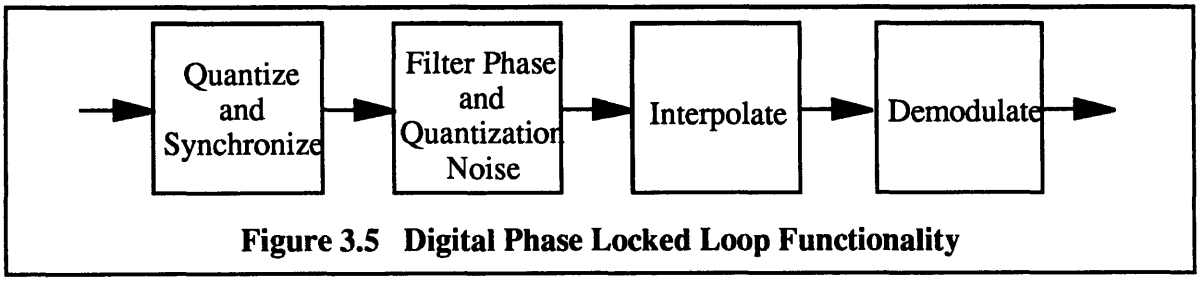

Figure 3.5 Digital Phase Locked Loop Functionality

Unlike the analog phase locked loop, synchronization (as well as quantization) is performed first rather than last with a digital phase locked loop architecture. Either the squared signal itself is synchronized before input to the DPLL, or the unsquared signal is compared to a synchronized, quantized bit, and the resulting error is measured using a synchronous counter, effectively quantizing any error measurement to a clock LSB. The

resulting quantization noise is the same either way, and the loop itself operates synchronously.

The loop itself is designed to filter phase noise, and potentially a certain amount of quantization noise as described below. If quantization noise is unfiltered by the loop, a certain amount of post filtering may be beneficial, as with the analog phase locked loop. The loop NCO effectively interpolates a phase measurement for each digital clock pulse. Hence measurements may be synchronously strobed from the readout on any clock cycle. Demodulation (carrier removal) may be accomplished directly, by strobing the digital output phase at the carrier rate or by subtracting digitally calculated nominal carrier frequency from the phase output. Data rate conversion is presumably accomplished through down sampling of the resulting phase measurements.

The digital phase locked loop may be implemented with memory by keeping track of many low frequency NCO bits. As with the analog phase locked loop, should the carrier "disappear" for some amount of time, the loop will continue to accumulate phase and resume lock when the signal returns, so long as the carrier frequency hasn't shifted substantially during the event.

Concerning the preprocessing of phase signals, certain tasks, such as periodic error removal, can be easily implemented as they were with the analog loop.

3.2.2 Limitations of the Digital Phase Locked Loop

The major limitation of the digital phase locked loop is resolution. Resolution limitations beyond dynamic tracking errors are the result of digital circuitry operating at a finite clock rate. Even if the input signal is not explicitly synchronized with a flip flop before phase detection, implicit quantization occurs when a digital phase detector cannot measure phase errors smaller than a clock pulse in length.

Independent of the phase detector resolution, a numerically controlled oscillator produces an output signal whose rising and falling edges occur at a clock edge, regardless of where in time the input signal's rising edges occur. Hence a perfect phase lock cannot be established to any signal if the zero crossings occur at times other than clock zero crossings, a condition that must be expected. In addition, with a simple counter NCO, resolution is limited by the number of bits in the NCO. Since a perfect lock is never established, the best results attainable occur if quantization noise is shaped, to the extent possible, by the loop filter, requiring that a special NCO and loop filter of some sort be used. Similar results are attained with subsequent filtering of the phase measurements reported by a simple loop, if the data is taken at a sufficiently high rate, and the combined quantization noise of the NCO and the phase detector is shaped away from the signal spectrum. A more complete discussion of the ultimate resolution which can be attained from a digital phase locked loop with synchronization induced quantization error (resulting either from the NCO or from the phase detector) follows from the Digital Clock Strobing Readout analysis in Chapters 4 and 5, and is offered in section 6.1.

In general, phase resolution of the digital loop is limited by the resolution of the digital clock which drives the circuitry, though this limitation may improve for certain signals, as will be demonstrated later. This resolution limitation is analogous to the quantization resolution inherent with the VCO frequency for the analog loop. The DPLL avoids quantization in this sense, since the demodulating signal is synchronous with the clock, but introduces quantization at the phase detection stage. This leads to the conclusion that in general, highly accurate phase measurement and digitization demand high frequencies, or accurate time measurements, in the readout circuit.

As with the analog phase locked loop, using digital phase locked loops requires that each signal have it's own phase locked loop which is either built or programmed especially for the particular signal.

3.3 Conclusions

Phase locked loop based readout implementations are popular and practical because of the fact that many of the functions of phase readout and digitization are accomplished at the same time with a small amount of hardware. Initial analysis indicates that resolution limitations occur as a result of quantization errors which are induced by either finite VCO or clock frequencies. The major disadvantage of phase locked loop based approaches is that each signal requires it's own, individually designed phase locked loop for demodulation and digitization. The fundamental advantage of an analog loop over an all digital loop is that a phase error null can be achieved, promising the *potential* for improved measurement accuracy.[1]

[1] Appendix C considers several hybrid phase locked loops which use an analog phase detector to achieve this same advantage.

Chapter 4: The Digital Clock Strobing Readout

The digital clock strobing readout, briefly discussed in the Introduction, offers several potential advantages over phase locked loop approaches. It is a flexible architecture, and does not require any analog components whatsoever. It does not require high frequency representations of the signal because it uses all of the bits of the system clock to obtain an accurate measurement of signal phase. In addition, it claims a maximum potential for shared resources, by using a single accurate and stable clock as a common metric for demodulating multiple instruments' signals.

4.1 Clock Strobing vs. Signal Strobing

A fundamental question of digitizing phase readout design asks how is the best way to measure the phase of a signal relative to an onboard clock. More generally speaking, the question of interest may be how to measure the *phase relationship* of two signals, regardless of whether one, both, or neither of the actual signals are purely sinusoidal (no phase variations with respect to real time).

A simple example demonstrates this quite clearly. Consider two perfect sinusoids that with a nominal frequency ratio of 4/3. Signal A goes through 4 cycles for every 3 cycles of Signal B.

$$f_a = \frac{4}{3} f_b$$

If we define a "phase relationship" as the total instantaneous phase of one waveform minus the total instantaneous phase of the other waveform (with a sign convention), then the phase change for a single beat cycle of the two waveforms in this example is $\pm(8\pi - 6\pi) = 2\pi$. Each time a cycle is "skipped" between signal A and signal B, signal B phase has advanced by 2π radians over signal A. Hence, we can define a relative phase term, $\Delta\phi$, as:

$$\Delta\phi = \phi_a(t) - \phi_b(t)$$

Signals A, B, and $\Delta\phi$ are pictured in figure 4.1, through a single beat cycle.

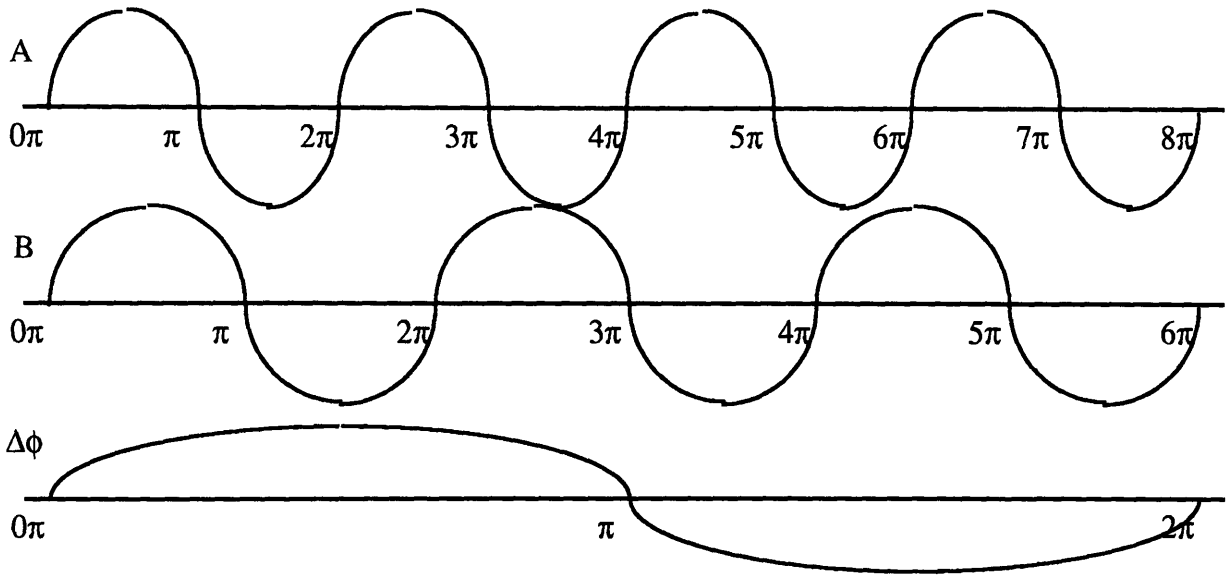

Relative Phase of Two Sinusoids

Figure 4.1 Phase Relationship of Sinusoids

In fact, if only two waveforms are considered, there is no way to tell which waveform is "moving" and which is stationary. The only real and meaningful measurement to be made is the phase relationship between the two waveforms. The simple point here is that there is only relative information to be retrieved.

In the case of squared signals, a phase difference is also definable. The phase difference is depicted in figure 4.2 below for squared up signal and clock frequencies which are *close* though not necessarily exactly the same.

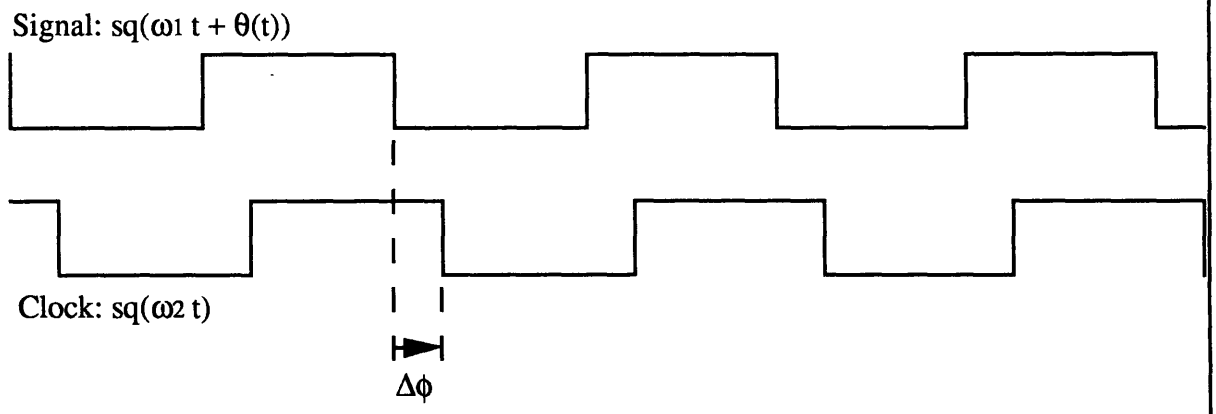

Figure 4.2 Phase Relationship of Square Waves

The case depicted here is analogous to the accelerometer problem. We call one waveform a "clock" and the other a "signal", and consider the clock to be a "reference". Clearly, measuring the phase relationship between the two waveforms does not require that one waveform be fixed in phase while the other is allowed to wander. A single variable is used to represent the time variation in phase rate between the two waveforms for ease of calculation, and since there is in fact only one physical parameter to be measured: the relative phase between the two waveforms. The sq function used to form the above waveforms can be defined as:

$$sq(\phi) = sgn(\cos(\phi))$$

To interpret exactly what is meant by the phase relationship of the two squared waveforms, we can define the instantaneous phase of each waveform to be a function of time equal to the argument of the sq function.

$$\phi_i(t) = \arg(sq(t)) = \omega t + \theta(t)$$

The instantaneous phase difference between the two waveforms can be defined as the difference between the arguments, as in the case of sinusoids.

$$\Delta\phi_i(t) = \phi_{i2}(t) - \phi_{i1}(t)$$

A problem inherent in these definitions is that the sq function has an infinitely multi-valued inverse function. The inverse cosine argument is periodically multi-valued, but more significantly, the inverse signum function is continuously multi-valued. As a result, knowing the value of the sq function tells us whether the argument, modulo 2π, is greater than or less than π. Hence, measuring the exact instantaneous phase of either the signal or the clock, relative to absolute time proves an impossible task at any point in time other than zero crossing times.

An approximate solution is apparent if we know that the instantaneous phase of a waveform (as a function of time) has an approximately constant rate of change. In this case, given the ability to continuously know absolute time, we could estimate the instantaneous phase function from the sq function at all times by interpolation between measurements. The accuracy of the estimation is based both on the reliability of the constant derivative approximation and the accuracy of the interpolation technique.

To determine how much information can be (theoretically) recovered from the sq function, we can use the fact that zero crossings of the sq waveform represent time samples of the instantaneous phase function at the times when the function is equal to 90 or 270 degrees. Typically, when we think of sampling a signal, we imagine measuring an arbitrary amplitude (within some range) of the signal at known times (typically periodically). In this case, however, the sq function effectively samples it's argument

whenever the argument reaches a certain value, but at arbitrary times. These values are evenly spaced in phase (every two pi) just as normal sampling intervals are evenly spaced in time. Regardless of the fact that the incoming instantaneous phase measurements may not arrive at a fixed rate (with respect to absolute time), we can still use the Nyquist theory to comment on the prospective ability to successfully interpolate phase values in between zero crossing times. If the phase variations are bandlimmited, so that

$$\Theta(f) = 0 \quad when \ f > f_{cutoff}$$

where

$$\Theta(f) = F(\theta(t))$$

it follows that so long as zero crossings occur more often than twice the highest frequency of phase variations, the phase function can be completely recovered for all times in between zero crossings as well.

This argument assumes the ability to tell absolute time with absolute accuracy, which in fact is not possible. Consider the general case where neither waveform is fixed in phase with respect to absolute time, perhaps neither meets the Nyquist condition with respect to absolute time, and no reference of real time is available other than the two waveforms. This is of no concern, if the measurement of interest is not the absolute phase of either signal, but the relative phase between them. To determine the appropriate condition for which the relative phase difference may be known at any point in time, simply consider either signal to represent absolute time, using any units. The choice is arbitrary. In this case, if the rate of change of the second waveform phase with respect to the first is such that less than half of a modulation cycle (with respect to the reference)

passes between edges, than the relative phase of the two signals may be known everywhere, (though the phase of neither waveform with respect to "real" time is known nowhere). The Nyquist criterion for the phase relationship is identical regardless of which waveform is considered a time reference. This remains true when a frequency offset exists in addition to phase fluctuations. Hence, it makes an equal amount of sense to calculate the difference between rising and falling edges of the two signals with respect to either waveform. One waveform must be used as a yard stick to measure changes in position of the other.

For this reason, there are two fundamentally equivalent ways to keep track of phase difference measurements between a signal and a clock. Either involves keeping track of the number of cycles (or fractional cycles) of "skip" between the two waveforms (which may or may not have the same nominal frequency), converting these cycle measurements into phase measurements , and subtracting out any constant frequency difference contribution to obtain the desired information, $\theta(t)$. Whichever waveform is to be used as a yard stick must be multiplied up in frequency (subdivided into half, quarter, eighth, . . . cycles) so that accurate measurements can be made of the relative position of the other carrier's zero crossings. A seemingly irrelevant choice is which waveform to use as a yardstick which is "strobed" by the other. Simply put, the question is should the signal strobe the clock or should the clock strobe the signal.

After some thought and analysis, several factors which might affect this choice have surfaced. One argument for clock strobing is to avoid the hassle of building an additional phase locked loop. For example, if a digital clock of several megahertz is available, and the phase of a signal of several kilohertz is to be measured, using the clock as a yard stick (strobing the clock) will yield close to a part in a thousand resolution per measurement even without resolution improvements from subsequent filtering.

A second reason why strobing a clock may be advantageous is that the clock is a

truer measurement of real time, potentially resulting in better measurements. While this argument seems to defy the previous analysis, the issue is a pragmatic one related to interpolating between zero crossings to establish fine resolution. Presumably, multiplying up either the clock or the signal requires the use of a phase locked loop, which contains a loop filter. As discussed previously, a finite order loop filter generates a phase error which depends upon the ratio between the carrier frequency and the frequency of applied phase modulation. This applied phase modulation is phase modulation with respect to real time, as determined by the loop VCO when a constant voltage is applied as input. Since a guidance system clock ideally runs at a fixed frequency with respect to absolute time, driving a phase locked loop from any of the clock bits results in static phase lock, with zero dynamic tracking error. Should a phase locked loop be synchronized to the signal, dynamic tracking errors result which are induced by the signal's changes with respect to real time. Hence clock strobing is advantageous when phase measurement accuracy is sought beyond the dynamic tracking limit of a realizable phase locked loop synchronized to the signal itself.

A final reason why clock strobing may be advantageous is that if a high resolution clock is available for one signal, it is available for every signal, i.e. the same clock can be used as a yardstick for every phase modulated signal which needs to be measured. Using signals as yard sticks requires an additional phase lock loop for each signal considered.

The primary motivation for avoiding clock strobing is that it substantially increases the complexity of subsequent digital processing. Since the measurements which result from clock strobing are not periodic, filtering may be more difficult. In addition, data rate conversion must involve more than up or down sampling of resulting measurements. Hence clock strobing potentially places a much heavier burden on digital signal processing circuitry.

4.2 Digital Clock Strobing Readout Functionality

As motivated by section 4.1, and the desire to implement an all digital phase readout and digitization architecture, the clock strobing readout, pictured again in figure 4.3 below, is alluring as a prospective candidate for system implementation.

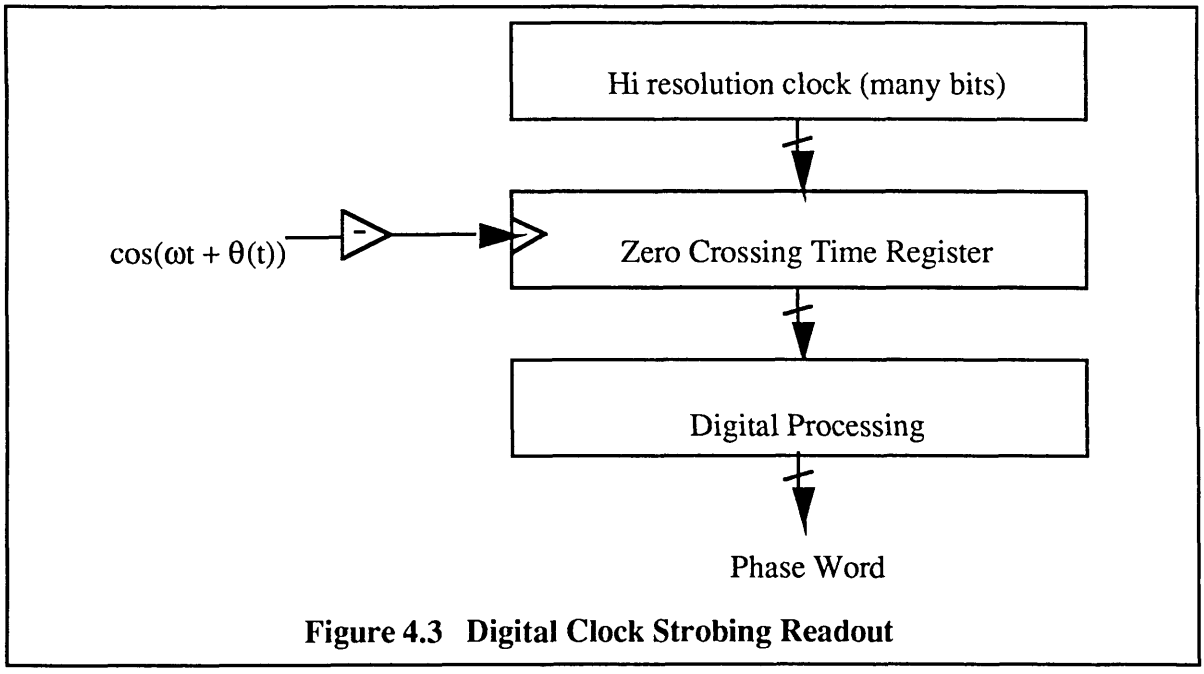

Figure 4.3 Digital Clock Strobing Readout

In terms of the phase readout and digitization functions defined previously, the clock strobing readout functionality is depicted in figure 4.4 below.

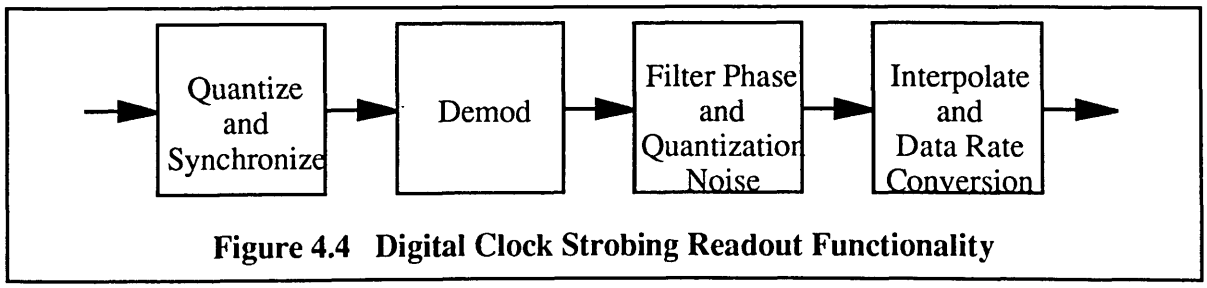

Figure 4.4 Digital Clock Strobing Readout Functionality

Synchronization takes place by quantizing the signal's zero crossing time to the time of

the most recent clock cycle. Phase noise removal, quantization noise removal, demodulation, and data rate conversion, along with preprocessing, are all tasks which must be implemented through the digital processing of clock phase words. Figure 4.4 depicts demodulation prior to filtering and data rate conversion, but this choice is somewhat arbitrary.

Because the clock itself exhibits memory, the readout exhibits memory as well. To the extent that the clock is uninterruptable, a circuit blackout period can result in an accurate measurement of carrier phase unless the carrier has shifted by more than 360 degrees in phase during the interruption period.

4.3 Resolution of the Digital Clock Strobing Readout

A fundamental limitation of the clock strobing readout architecture is the number of bits resolution of the clock itself. This statement directly parallels the conclusion that a digitizing phase locked loop is limited in measurement resolution by the number of bits in the feedback counter which is strobed. Correspondingly, define

$$b_o = \log_2 \frac{f_c}{f_n} + 1$$

where

$$b \equiv \textit{bits of phase resolution (per cycle)},$$

and

$$b_o \equiv \textit{resolution of clock relative to carrier}.$$

The 1 bit is added to reflect the highest clock bit has information as well, so the effective resolution of the clock is half a clock bit.

Under certain conditions, the Digital Clock Strobing Readout will not surpass clock resolution, b_o,. It doesn't follow, however, that because the resolution of the clock is limited to some number of bits, that phase measurements can't be improved under any conditions. Certainly, phase measurements are bound by the accuracy of the clock, (to the extent that it loses or gains time each cycle), but to the extent that quantization noise is removed from the resulting stream of measurements by subsequent filtering, it seems likely that better resolution may be obtained. In fact, the accuracy of a guidance system clock may be far better than it's resolution, suggesting that filtering of phase measurements could result in reliable measurement improvements. This notion is analyzed and tested throughout the remainder of section 4.3.

4.3.1 DC Resolution Limitations

Specifically, consider the digital clock strobing readout with a clock frequency precisely equal to 2^n times the nominal carrier frequency:

$$f_c = 2^n \times f_n$$

where

$$f_c \equiv highest\ clock\ frequency$$

and

$$f_n \equiv nominal\ carrier\ frequency.$$

Now, if the instrument, an accelerometer, were to enter a zero g environment, it's output would exhibit no phase variations, producing a pure carrier frequency. At this point in time, and for all future times before acceleration (phase) variations begin to occur, consider the accuracy which the instrument's phase (and the corresponding instrument velocity) can be measured.

The problem inherent in this particular situation is results from the fact that the clock and carrier are perfectly synchronized; they are in phase lock with one another. Hence no amount of measurement filtering will improve reported measurement resolution beyond the resolution of a single measurement. Every time a signal zero crossing occurs, it will fall exactly in the same place between clock bit changes as it has for each previous cycle, and for each following cycle (while the 0 g acceleration field is maintained). The situation is depicted in figure 4.5 below:

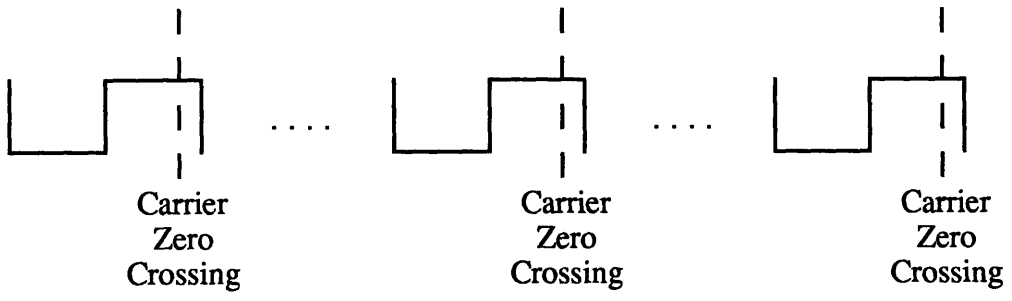

Figure 4.5 Phase Lock of Clock Synchronized Carrier

In this case, the maximum phase resolution that can be achieved, in bits, is clearly

$$b_o = n+1$$

i.e. if the clock is twice as fast as the carrier (n=1), a zero crossing time can be resolved to within 1 quarter of a carrier cycle (if the clock has a 50% duty cycle), represented by 2 bits of phase resolution per cycle. The minimum accuracy corresponding to 1/4 cycle resolution (2 bits) is therefore equal to 1/8 of a cycle if the reported measurement for a given interval corresponds to the middle of the interval (an extra bit of resolution is reported to minimized maximum error). Both the resolution, reported measurement, and accuracy for the previous example are indicated in the figure 4.6 below.

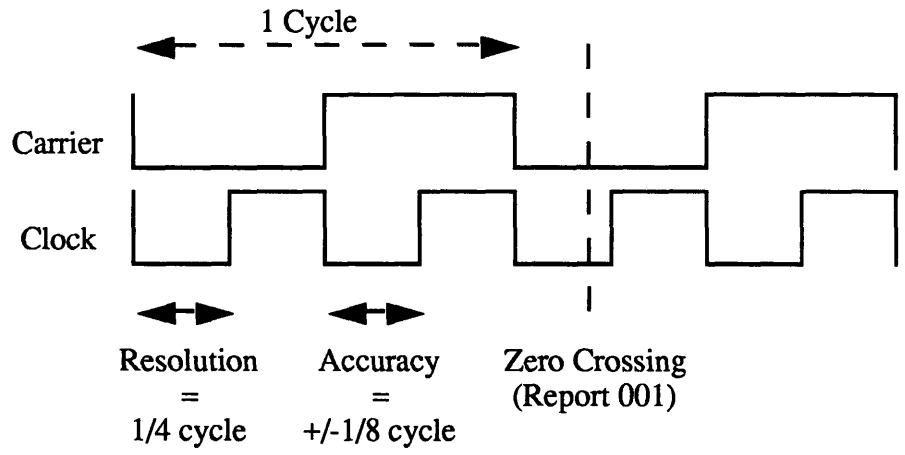

Figure 4.6 Phase Resolution and Accuracy Example

More generally speaking, a similar resolution limitation should exist whenever the clock to carrier frequency ratio is an integer. In this case, if

$$f_c = I \times f_n$$

where

$$I \equiv \textit{any integer}$$

then the resolution limitation equation takes the more general form

$$b_n = \log_2(I) + 1$$

which reduces to n+1 when I is set equal to 2^n.

More generally, a similar resolution limitation should exist whenever the ratio between the clock frequency and the carrier nominal frequency is a rational number. If we characterize the frequency relationship as:

$$N \times f_c = I \times f_n$$

where

$$N \equiv \textit{any integer}$$

and I is defined as a before, then an analogous situation exists to the previous example, with the exception that phase lock occurs over N carrier cycles instead of one carrier cycle. To establish the resolution limitation equation, consider a frequency ratio of 5/2 ($N = 2$, $I = 5$). The carrier cycles for this situation are depicted figure 4.7 below:

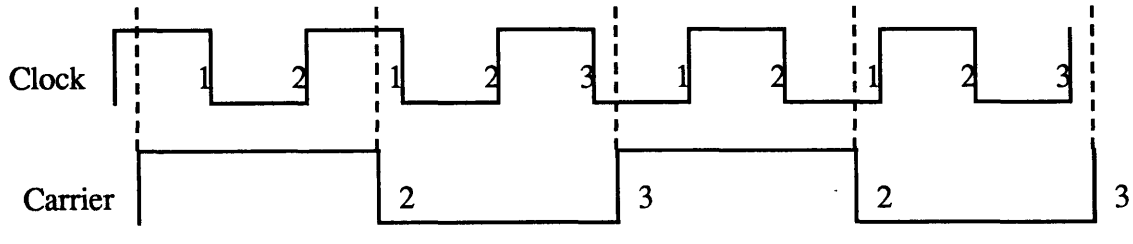

Figure 4.7 Rational Frequency Ratio of 5/2

The numbers adjacent to the clock cycles correspond to a phase count, accumulating to the numbers measured at each carrier zero crossing. Because $N = 2$, the resulting unfiltered phase measurement pattern, 2, 3, 2, 3 . . ., repeats every two half cycles. Hence any filtering of the raw measurements will effectively see only two measurements.

To see the resolution limitation resulting from the two waveforms above, consider how much the carrier would need to slide in phase, either forward or backward, before the resulting measurements could be distinguished from the original. A full 1/8 of a clock cycle advance or retard in phase would result in a different stream of output measurements, indicating that the resolution achievable in this situation is 1/4 of a clock cycle, corresponding to (2/5)*(1/4)=(1/10) of a carrier cycle. This is the same resolution which results from a clock frequency 5 times the carrier frequency (the extra factor of two resulting from the 50% duty cycle, once again). Based on this example, and some extrapolation, it can be seen that the phase resolution limitation equation with a rational frequency ratio as defined still takes the form

$$b = \log_2(I) + 1.$$

In simple terms, when the clock beats against the carrier over N carrier half cycles, the resulting measurement resolution limit should be the same as what could be achieved using a clock frequency N times as fast as the actual clock. I will refer to this property as the "N effect."

4.3.2 N Points and I Points

Analysis of the DC phase resolution limitations inherent to the clock strobing readout

offered in section 4.3.1 indicates that for any given clock frequency, an infinite, set of non-adjacent carrier frequencies exists such that DC phase resolution at one of these frequencies will be limited by the DC phase resolution limitation equation. This set of frequencies comprises two sets: the set of frequencies which are equal to the clock frequency divided by an integer, and the set of frequencies which are equal to the clock frequency times any other rational number. I will refer to these frequencies as I Point and N Point frequencies, respectively. "I Point" frequencies may be calculated, with respect to a certain clock frequency as

$$f_I(I, f_c) = \frac{f_c}{I}.$$

Likewise, N point frequencies may be calculated for a certain clock frequency and I value, as

$$f_N(f_c, I, N) = \frac{N f_c}{I}.$$

From a design perspective, however, it's more useful to find N points within a band of closely related frequencies. The N points frequencies in the neighborhood of a particular I point frequency, may be calculated as

$$f_N(f_c, I_{close}, N) = \frac{f_c}{I_{close} + 1/N}$$

where I_{close} corresponds to the value of I at the nearest I Point frequency. DC phase resolution at these frequencies with be limited according to the resolution limitation equation for N and I, with

$$I = N \times I_{close}.$$

resulting in a proximity resolution limitation equation of

$$b = \log_2(I_{close} \times N) + 1$$

4.3.3 Response to a Phase Ramp

The analysis from section 4.3.2 applies to phase ramps as well as constant phase inputs. Considering an accelerometer, for example, the same reasoning which indicates that resolution limitations could exist for the accelerometer in a 0 g field can also be used for a 1, 2 or any g acceleration profile.

In other words, if the readout attempts to digitize a phase ramp defined by

$$f_r = f_n + \frac{d\theta(t)}{dt} \times \frac{1}{2\pi}$$

where

$$\frac{d\theta(t)}{dt} = K \equiv some\ constant$$

then the phase resolution limitation is formed for this phase ramp frequency, f_r, instead of the nominal frequency used in section 4.3.2, i.e.

$$b = \log_2(I) + 1$$

where in this case, I is defined by

$$f_c = I \times f_r.$$

Hence, if the DC phase and phase ramp resolution limitation equations are to be utilized to estimate a bound on performance, the full range of frequencies for f_r, determined by bounds on $d\theta/dt$, must considered as potentially resolution limiting frequencies.

4.3.4 Performance Bounds for SFIR Instrument Readout

A specific example, using the SFIR instrument's phase characteristics, demonstrates the formation of DC and phase ramp resolution limitations over a range of possible fixed frequency outputs.

The SFIR (Specific Force Integrating Receiver) produces a nominal carrier frequency of 6.4 kHz.[1] This signal is the output of a resolver which is driven by a bit of the system clock itself. The highest clock frequency is 2^{10} (1024) times higher than 6400 Hz. In this case, the nominal frequency itself, as a simple fraction of the digital clock frequency, establishes a DC phase resolution bound for the instrument at 11 bits per cycle.

The two adjacent frequencies with associated I point phase resolution limitations are at (1024/1025) and (1024/1033) times 6.4 kHz, at approximately 6.4 kHz +/- 6 Hz. With an instrument scale factor of 1 radian/second/g, and a presumed maximum acceleration magnitude of 10g, the instrument's instantaneous frequency may deviate by only 10 radians/second, less than 2 Hz. Hence, neither of the resolution boundaries at the adjacent frequencies will affect performance.

Calculation of adjacent I point frequencies seems unnecessary, to the extent that a

[1]The SFIR actually uses a multi-speed resolver, with 1 and n speed. A quadrature phase signal is also present for both resolvers.

bound on DC performance at either of these frequencies would differ only slightly from the DC performance bound at the nominal frequency. What is interesting, however, is that even with a fairly large clock to carrier frequency ratio, a significant band of 6 Hz, substantially higher than the instrument's frequency deviations (4 Hz) exists. Were the nominal frequency to lie, for example, halfway between I points at 6.4 kHz plus or minus 3 Hz, a factor of two, or 1 bit improvement in the DC resolution limitation bound would result. Were the nominal frequency to lie at 6.4 kHz plus or minus about 2 Hz, a factor of three improvement in the DC resolution bound would result. No further improvements can be guaranteed because N point frequencies for larger N get spaced more tightly in the spectrum. Hence the DC phase resolution of the SFIR cannot surpass 12.6 bits, even if the instrument is operating at a shifted nominal frequency (non-zero acceleration) moving it away from the I point at 0g.

4.3.5 Effect of Frequency Deviation Ratio

The SFIR analysis suggests that some general bounds on the DC and phase ramp resolution limitation exist, given a particular frequency deviation ratio of a phase modulated signal. Clearly,

$$\lim_{f_c \to \infty} b = b_o = \log_2\left[\frac{f_c}{f_n}\right] + 1$$

independent of whether the frequency ratio is an integer or rational number. This is true for signals with non-zero frequency, which are the only signals of any interest (if the frequency and phase never changed, we wouldn't be very interested in measuring it). The limit is due to the fact that as the frequency ratio increases, the set of frequencies which are simple fractions of the clock frequency move closer together, so eventually even the

slightest frequency deviation range would limit DC resolution to the resolution of a single measurement.

A more useful bound on the DC resolution limit can be determined directly from a frequency deviation ratio, defined as

$$FDR = \frac{\Delta f}{f_n}$$

where

$$\Delta f \equiv 2\ max\left(\frac{d\theta(t)}{dt}\right).$$

In terms of this ratio, we can say that for $FDR \times f_c \geq f_n$, $b \approx b_o$, where the approximation occurs because an I Point frequency may or may not occur at the nominal carrier frequency. For smaller frequency deviation ratios, optimal positioning of the center frequency with respect to the clock frequency results in a maximum DC phase resolution limit of

$$b_m = \log_2(1/FDR) + 1,$$

Resolution limitation improvement can only occur if the band of instantaneous frequencies occupied by a signal contains no simple fractions of the clock. Likewise, a smaller frequency deviation ratio permits the band of instantaneous frequencies to squeeze between an I point frequency and increasingly close N point frequencies.

4.3.6 Effect of Phase Noise on DC Phase Resolution Limits

The analysis of sections 4.3.1 through 4.3.5 has assumed the absence of noise. With the introduction of phase noise, DC phase resolution limits may improve, so long as noise does not corrupt the signal itself. Consider the introduction of a phase noise term, $\theta_n(t)$, to a phase modulated signal, $s(t)$, such that

$$s(t) = \cos(2\pi f_n t + \theta(t) + \theta_n(t))$$

where $\theta_n(t)$ is characterized by an mean value of 0 and variance, σ^2. Given the condition

$$\frac{2\pi\sigma^2}{f_n} > \frac{1}{4N \times f_c}$$

the DC resolution limitation should cease to be a limit. The reasoning behind this statement stems from the fact that zero mean phase noise which deviates beyond the phase resolution of the clock ought to produce a filtered average closer to the signal phase. For noise which stays within the threshold of a half an effective clock cycle (with higher values of N resulting in higher effective clock frequencies), no amount of averaging can create this result. Hence the second term of the inequality is set equal to half of this threshold, as the noise exhibits both the positive and negative phase variations.

The implications of noise threshold condition are threefold. First, raising the clock frequency (shrinking the period) to within the phase noise threshold may offer performance substantially higher than what was obtained at slightly lower clock frequencies. Second, improving clock resolution beyond the noise detection threshold will offer rapidly diminishing returns. Finally, the controlled addition of phase noise

prior to detection may offer improved DC performance without the use of a high frequency clock.

4.3.7 Dynamic Performance Characteristics

The usefulness of the DC resolution limitation in predicting resolution attained in a dynamic phase environment is questionable. Certainly, it has merit for modulations which vary extremely slowly relative to the carrier frequency, i.e.

$$\Theta(f) \approx 0 \, for \, f \leq f_{cutoff}$$

and

$$f_{cutoff} <<< f_n$$

where the extra "<" is included for emphasis. To the extent that significant changes of phase take place over a finite number of cycles, the effective number of zero crossings used to estimate phase lessons, dramatically reducing the "N effect" described previously, especially for large values of N which require filtering over many carrier half cycles.

To approximate the affect of dynamic phase variations on potential resolution limitations, consider the amount of averaging required to effectively improve quantization resolution for a given value of N. Since averaging must occur over an entire beat cycle to improve resolution to the DC limit, it should be expected that for phase modulations which induce a substantial phase change during a beat cycle will not see the benefit of the N effect. A good approximation of a "substantial phase change" is any change of phase more than the resolution sought by averaging over the N cycles.

Considering the case where N=2, for example, if a modulation induced phase change of more than 1/4 of a clock cycle occurs during the time it takes for two carrier half cycles to complete, than resolution should drop from the value predicted by the resolution limitation equation.

4.3.8 Other Resolution Metrics

It is not clear that worst case phase resolution is the parameter of interest for a particular readout system. In fact, a better metric for many applications may be some sort of average resolution, though the two measurements are apt to be highly correlated.

In the case of inertial instruments, both metrics ought to be given consideration, as well as a third and possibly most important parameter: integrated error. Consider the accelerometer problem, where phase measurements are indicative of velocity. If the ultimate calculation of interest is position, which to first order will be proportional to integrated velocity, or phase measurements (assuming a fairly linear device), then integral phase error is the most important parameter. Unfortunately, very little can be said about integrated phase error due to instrument readout limitations unless a great deal can be said about the specific phase modulated signal.

4.4 Digital Processing of Zero Crossing Times

While section 4.3 addresses the phase resolution limitations of the Digital Clock Strobing Readout, it does not address in what manner carrier zero crossing times ought to be processed to yield the ultimate phase relationship measurements

Several approaches for digitally processing zero crossing times to produce desired filtered phase information have been considered. One approach involves using the zero

crossing times to calculate period measurements, converting these measurements to frequency measurements, subtracting a nominal frequency, and integrating these frequency differences to obtain desired phase information. A second approach conceived is to utilize a digital phase locked loop recursive algorithm, whereby each zero crossing time is predicted, measured, and the resulting error is used to adjust the phase and frequency estimates for the next prediction. A third method is to directly generate a phase estimate from the zero-crossing times by subtracting off an accumulated nominal frequency word and dividing by a nominal period, and then filtering the phase measurements directly.

4.4.1 Period Readout Method

The driving concept behind this method is than each zero crossing corresponds to a new period measurement. If successive zero crossings are subtracted to generate period words, then these period measurements may be converted to frequency, demodulated by subtracting off a carrier component, and integrated to yield accumulated relative phase data. The approach is depicted in figure 4.8 below.

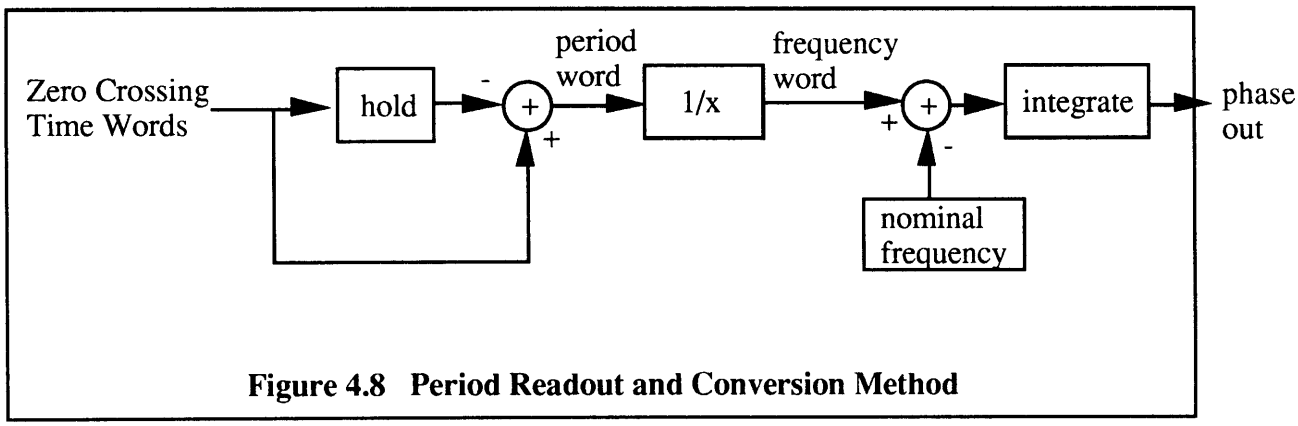

Figure 4.8 Period Readout and Conversion Method

Using this conversion method, digital filtering could be done in one of four

locations: at the zero crossing word input, the period word output, the frequency word output, or the phase word output. Notice that the interpolation and data rate conversion is not necessarily performed in this processing implementation, the integrator could perform such a task, effectively performing a linear interpolation of phase corresponding to the previous carrier half cycle.

Without getting into great detail, the fundamental problem with this approach is that phase errors will undoubtedly accumulate. We must assume that the nominal frequency estimate is correct, and enough bits are present to prevent any substantial error from accumulating over critical time periods If this were not the case, no readout methodology will be adequate. Even with a perfect carrier, however, and the assumption that over time, no error accumulates in the period word measurements themselves, the 1/x function threatens to become a serious source of error, in part because of approximation, but more so because the benefit of the period values average integrity is lost, i.e. if

$$\frac{x_1 + x_2}{2} = x_0$$

then

$$\frac{1}{x_1} + \frac{1}{x_2} \neq \frac{2}{x_0}$$

unless the two values are equal. In the case the x's correspond to period measurements which do indeed carry meaning in their average value, and hence the errors associated with each measurement may accumulate instead of averaging away. The result of using this method is to increase the resolution demand for each individual measurement to a value which requires a multiple gigaHertz frequency clock.

4.4.2 PLL algorithm

A second approach considered for the processing of carrier zero crossings is a digital phase locked loop algorithm. The essence of this approach is to predict when each zero crossing will arrive, and generate an error term based on the actual measurement which affects the next prediction, just as a digital phase locked loop would do. The predicted zero crossing times can then be used to extrapolate phase at any arbitrary time. A potential implementation of this methodology is depicted in figure 4.9 below.

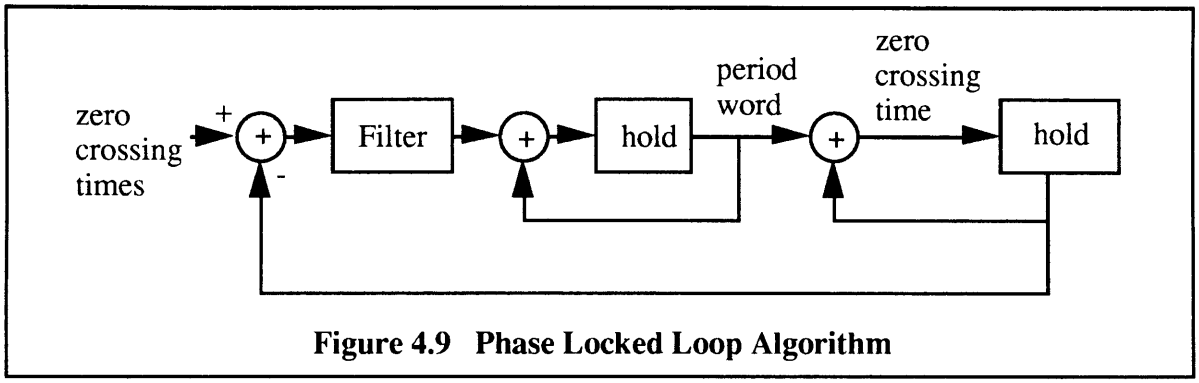

Figure 4.9 Phase Locked Loop Algorithm

The idea here is that all of the functionality of a digital phase locked loop can be implemented, and more. This is true because the time of each zero crossing is known to within the resolution of the clock. Consequently, if a zero crossing time is predicted (somehow, anyhow) a phase error between the predicted value and actual value can be measured to within the clock's resolution, just as a digital phase locked loop would do. The advantage in this case is that it is not necessary to generate a physical demodulating bit. This allows considerably more flexibility in the way the loop filters errors and predicts a zero crossing time. In addition, the loop can effectively filter quantization noise as well as phase noise, since the only source of quantization noise will be synchronization noise. Zero crossing predictions can contain an arbitrary number of bits.

To generate relative phase measurements from the filtered zero crossing times (predictions), several options are available. Bandlimited interpolation of arbitrary order could be used to generate phase values at times which correspond to clock times, and nominal phase values equal to a running integral of nominal frequency could be subtracted off at these times to yield relative phase numbers. Alternatively, the period values corresponding to the filtered zero-crossing times could be input into the period conversion approach discussed previously, with the advantage that some of the averaging which could not occur before (filtering of synchronization induced quantization errors) would have already been accomplished presumably resulting in smaller integrated errors. Of the demodulation methods here, the first will most likely result in higher quality data and less integrated error, but the second approach may be much less computation intensive.

An alternate approach to interpolation, more closely related to the operation of a typical digital phase locked loop, is to use an NCO which generates a carrier phase estimate for every digital clock tick. These values could be used to estimate zero crossing times for error detection (subtraction) as well. Thought of in this way, it becomes clear that any digital phase locked loop (with synchronization induced quantization error) which is physically realizable can be implemented from the zero crossing time words. The clock strobing readout, in this context, is simply playing the role of phase detector for the loop, and providing memory so that the readout circuitry can easily figure out when and how many missing pulses occur in a given signal outage.

The only weakness with this implementation is that only recursive filtering is implemented, as is the case with phase locked loops in general, which are inherently recursive. The trouble here is that dynamic performance is bound to suffer because the filtering is causal. Avoiding dynamic signal induced phase errors (which may integrate into velocity errors) requires non-causal FIR filtering.

4.4.3 Direct Phase Calculation

A third conceivable approach to generating phase words from zero crossing times is what I call direct conversion. This approach may be the simplest and most elegant method. The method proposes to subtract from each zero crossing time word the time that a nominal frequency carrier would have reached the same total phase. The difference is a scaled unfiltered version of the desired phase measurement. This approach is depicted in figure 4.10 below.

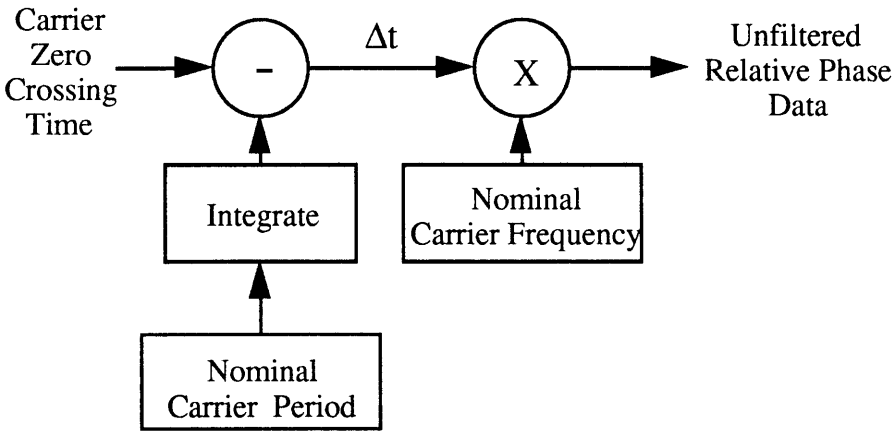

Figure 4.10 Direct Conversion Method

The relative time or phase data can be filtered, although I have depicted the actual phase measurements being filtered. The counterintuitive notion about this approach is that demodulation takes place using period values, rather than frequency values, from one perspective. From another perspective, the demodulation is taking place in terms of phase, rather than frequency, but in clock phase instead of carrier phase before conversion takes place (multiplication by nominal frequency) Either way of thinking about it, the effective result is identical to what was accomplished by the period

conversion method, with the exception that phase average integrity is maintained, resulting in no integrated errors other than those attributable to a bad value of nominal frequency. This simple and accurate means of demodulating leaves plenty of options open for filtering, as the resulting relative phase measurements can be filtered recursively or FIR filtered.

Overall, a digital phase locked loop processing scheme makes the most sense if a lot of things are to be accomplished with a small amount of dedicated hardware, such as a special NCO and recursive filter of some sort, because the loop can effectively demodulate, interpolate and filter at the same time. If high power general purpose digital signal processing is available, however, the direct conversion approach is most desirable because it handles demodulation and lets DSP hardware do the serious number crunching to obtain near optimal performance.

4.5 Conclusions

The Digital Clock Strobing Readout architecture effectively synchronizes data to the clock, and provides memory through it's own state. As a result, the values strobed from the Readout are analogous to the values strobed from a digitizing phase locked loop, but filtering, interpolation, and dynamic tracking issues can be handled independently. The Digital Clock Strobing Readout, incorporating demodulation by direct conversion, and combined with digital signal processing power, should be able to outperform any digital phase locked loop implementation (with synchronization induced quantization of error detection)

The ultimate resolution which can be achieved by the readout will still fall within the limits imposed by the frequency relationship between the clock and carrier frequencies, as derived in section 4.3. But for the right combination of frequencies and

sufficiently slow phase variations, readout resolution should improve substantially beyond the clock's resolution.

The danger which could mask the performance of the DCSR is if an approach such as the period readout method described is implemented somewhere in the processing, effectively throwing out the information in the average phase measurements which allows measurement accuracy finer than the clock's resolution to be obtained.

Chapter 5: Simulation of Digital Clock Strobing Readout

The digital clock strobing readout was simulated using a spreadsheet simulation template. The context of the simulation is an accelerometer readout, with an output frequency and phase proportional to acceleration and velocity, respectively. This simulation, explained in detail in Appendix A, permits testing of resulting phase measurement accuracy as certain parameters, such as nominal carrier frequency, clock frequency, instrument scale factor, and filtering coefficients are varied. The output parameter of interest from the simulation is bits of resolution, corresponding to number of correct bits. Both average and worst case measurement resolution is computed for a particular set of simulation parameters.

The simulation implements the direct phase conversion method described in section 4.4.3. The resulting measurements are asynchronous with respect to time. Further processing of the measurements would be required in a real system implementation, though adequate interpolation techniques should not affect the quality of the data. This statement is based on the assumption that carrier zero crossings, arriving at approximately twice the carrier frequency, occur much faster than the signal itself changes, permitting bandlimited interpolation.

5.1 DC Performance Characteristics

Several sets of simulations were performed using constant or zero acceleration profiles (corresponding to phase ramps), as well as variable clock and carrier frequencies. These simulations were performed to verify the DC resolution limitation equations discussed in Chapter 4.

5.1.1 Effect of Carrier Frequency on DC Resolution

The response of the system to a set of constant acceleration profiles (and corresponding phase ramps) was simulated, and both the minimum and average bits of resolution attained, after filtering, were logged. For this particular simulation, a clock frequency 2^{10} times the carrier frequency was used, and filtering was with a 21 point rectangular finite FIR (finite impulse response) filter (a moving average). A scale factor of 1 rad/sec/g (about .16 Hz/g) was used. The simulations ran over 100 carrier cycles. The resolution measurements from this set of simulations are depicted in figure 5.1, with end effects removed from resolution calculations.

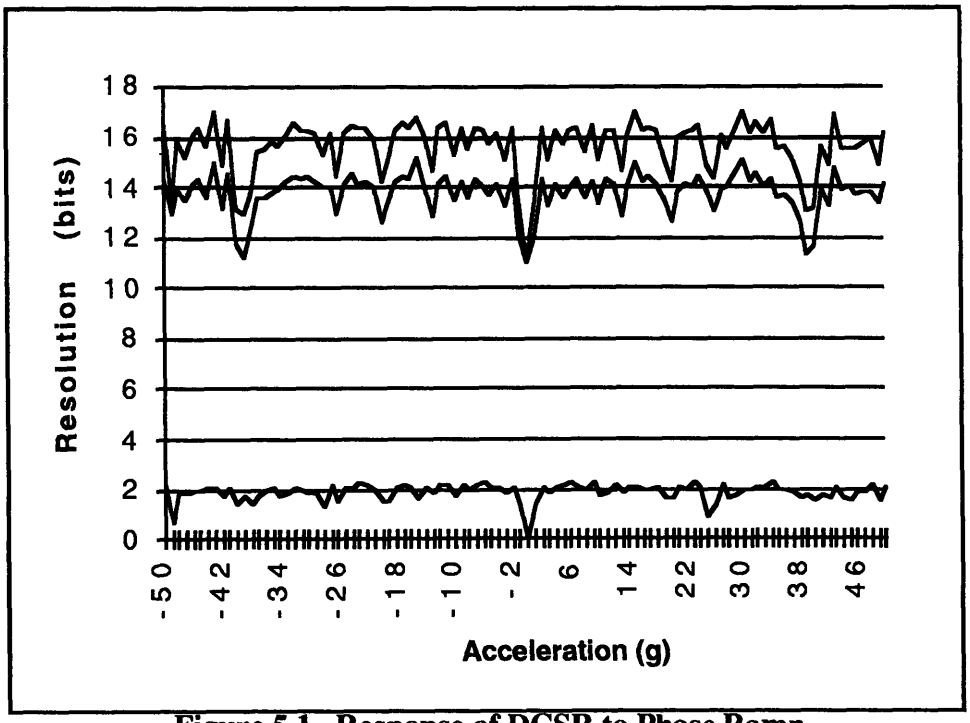

Figure 5.1 Response of DCSR to Phase Ramp

The top line is average resolution. The second line is worst case resolution. The bottom line represents the difference between best and worst case resolution.

A first glance at the simulation results confirms some intuitive notions. Note that the resolution achieved at 0 g, when the carrier and clock are in perfect lock, is exactly 11 bits, both in average and worst case value. As explained previously, this resolution corresponds to the clock resolution itself, as predicted in section 4.3.1. In this particular example, n=10 and b=10+1=11. Likewise, it is intuitive that the worst case and average resolution are identical in this case, since the resolution of every measurement is exactly 11 bits. Observing the near symmetry of the resolution measurements confirms a second intuitive notion: that DC phase resolution for a given phase ramp ought to be symmetric in the neighborhood of a perfect phase lock, since N points occur on either side of an I point at approximately equal and opposite frequency differences.

Notice also that the minimum resolution achieved is greater than the 11 bit

minimum in all cases, except in the neighborhood at -39 and 39 g. These accelerations correspond to $6400Hz \pm 39 \times (.16Hz) = 6400 \pm 6.3Hz$. Given a clock frequency equal to 1024 times the frequency of the carrier, we would expect the resolution minima to occur at 1024/1025 and 1024/1023 times the carrier frequency, equal to 6393.8 Hz and 6406.3 Hz respectively, or $6400 \pm 6.3Hz$. Other notable minima occur at ± 13 ± 20, and ± 26 Hz of 6.4 kHz , each with approximate values between 12.5 and 13 bits. These points nicely correspond to values spaced approximately 1/3, 1/2, and 2/3 between the two phase lock points. Correspondingly, we should see factors of 2 and 3 higher resolution measurements (N=2 or N=3) at these values than at 0 or 39 g, corresponding to 1 and 1.6 bits. While the simulations indicate an additional 1.5 to 2 bits at these accelerations, the acceleration values are not at the exact minima points.

Interestingly enough, this set of simulations showed a remarkably consistent difference between average and minimum resolutions achieved, hovering near two bits. This result seems to indicate that if average resolution is the metric of interest, the DC minimum resolution boundaries established in section 4.3 may be used, along with a fudge factor of 1.5 to 2 bits, as a rule of thumb for predicting average DC resolution.

Since DC resolution limitation is determined by the clock/carrier frequency relationship, it follows that the results of varying the carrier frequency over a particular band may be duplicated by varying the clock frequency over a related band of frequencies. For example, the previous set of simulations varied DC acceleration inputs over a range from -50 to 50 g, corresponding to radial frequency shifts of -50 to 50 rad/sec. An equivalent band of clock frequencies, defined by upper and lower bounds of f_{c+} and f_{c-}, is determined from upper and lower carrier frequency bounds f_{n+} and f_{n-} by preserving the frequency ratios

$$\frac{f_n}{f_{c+}} = \frac{f_{n-}}{f_c} \text{ and } \frac{f_n}{f_{c-}} = \frac{f_{n+}}{f_c}.$$

Hence a set of simulations with clock rates defined by these ratios using the acceleration values from the previous experiment produces an identical plot, shown in figure 5.2 below[1].

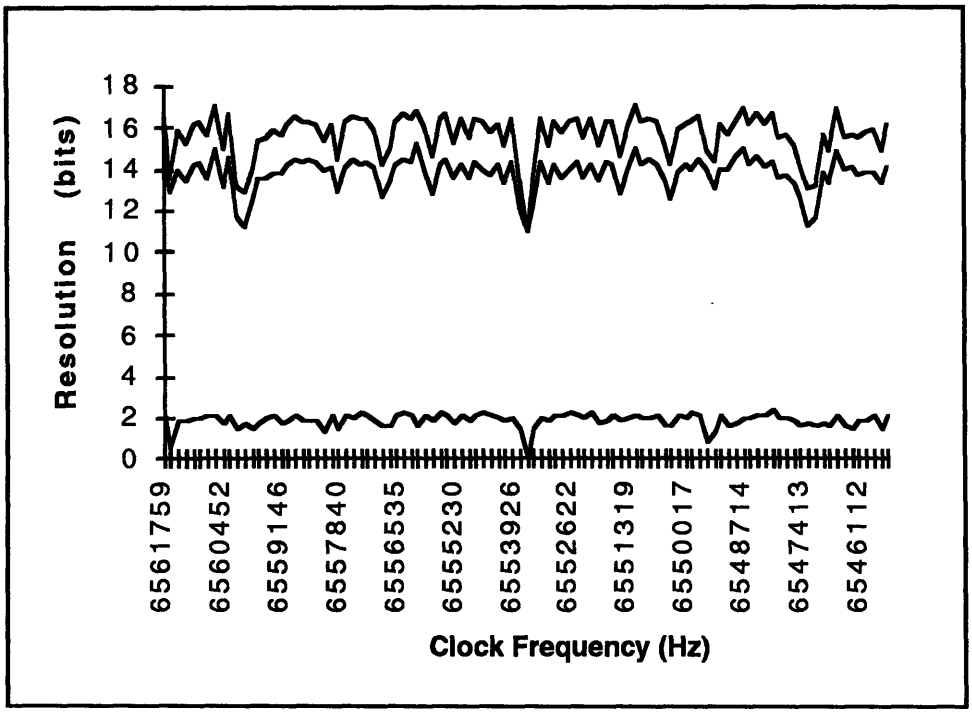

Figure 5.2 Clock Frequency Effects on Resolution

Overall, the DC phase ramp and clock frequency varied simulations demonstrate clearly how DC resolution improvements can be achieved beyond a clock's resolution within certain frequency bands isolated from I points or N points where N is small. Such an area is found in this simulation between the values of 6400 and 6413 Hz. For DC frequencies within this range, resolution improvements of 2 to 3 bits can be expected with a small amount of filtering.

[1] In fact, the same exact frequency ratios were used at all points by setting each clock frequency value simulated equal to nf*(6400*2^10)/(nf+(M17*sf)) where M represents the column of acceleration values simulated in the previously described experiment.

5.1.2 Filtering Effects on DC Performance

A similar set of simulations was performed to the ones described above, to demonstrate the affects of improved filtering upon resolution, and obtain a cleaner picture of DC resolution limitations, particularly for N larger than 3. Using the same parameters as before, this set of simulations ran over 1000 carrier cycles and filtered over 201 cycles. Readout responses to accelerations ranging from 0 to 40 g in increments of .2g were computed, yielding the chart in figure 5.3.

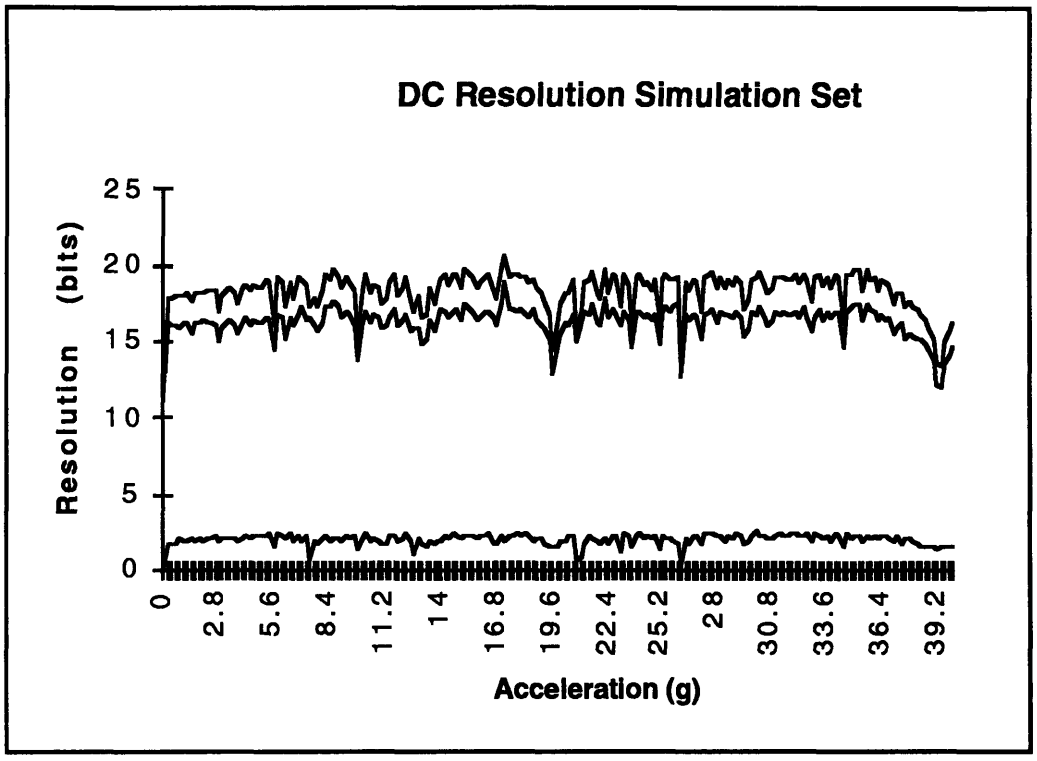

Figure 5.3 DC Resolution With Additional Filtering

With the additional averaging, and smaller acceleration increments, a more detailed picture of DC phase ramp performance is apparent. Many additional resolution minima are increasingly visible, though few points are distinct. The results here are somewhat deceiving. It appears at a first glance that the resolution at N Points has

improved beyond the limit as a result of additional filtering. This is not actually the case, as will be shown in section 5.1.3. In fact, despite finer granularity in the acceleration vector (x axis), more filtering has a dramatic improvement on resolution near N points, masking the location of the N points between simulated acceleration increments. This property is discussed more in section 5.1.4.

To get a better feel for how the additional filtering improved resolution, consider the average over all frequencies of minimum and average resolution, depicted in table 5.1, for the two simulation sets.

Table 5.1 Filtering Effects Across Carrier Spectrum

	Min.Min.	Avg. Min.	Max Min.	Min. Avg.	Avg. Avg.	Max Avg
F=21	11	13.765	15.189	11	15.635	17.079
F=201	11	16.263	18.869	11	18.292	20.711

Using these two filter lengths for data points suggests an average DC resolution improvement of about 2.5 bits per factor of 10 in filter length, with substantially higher improvement of about 3.5 bits in maximum resolution (at particular frequencies).

One would expect that with filtering across F points, resolution at N points as high as N=200 should obtain the values predicted by the resolution limitation equation, corresponding to an additional 7.6 bits of minimum resolution beyond the value achieved at I_{close}, or 18.6 bits, quite close to the 18.9 bits indicated in the table above. Likewise, an additional 4.3 bits of minimum resolution should occur at some frequencies for f=21, quite close to the 15.2 bits indicated above. Hence, in appropriate bands of frequencies, an additional 3.3 bits of resolution (factor of 10 in resolution) may be achieved for each factor of 10 increase in filter length.

5.1.3 Exact Resolution at N Points

Since the previous simulation sets did not clearly show resolution limitations at "N points", several simulations were performed at accelerations precisely equal to those where resulting frequencies should produce local resolution minima. Specifically, since "N points" in the neighborhood of a particular I occur at frequencies

$$f = \frac{f_c}{I_{close} - 1/N}$$

I chose accelerations which correspond to N points in the vicinity of 6.4 kHz by setting

$$a = \left(\frac{f_c}{1024 - 1/N} - 6400Hz \right) \div sf$$

yielding the values in table 5.2.

Table 5.2 Resolution at N Points - Part 1

N	Acceleration	Min Bits	Avg.	Difference	Predicted
1	0	11	11.3607732	0.36077319	11
2	19.64452956	11.9928402	12.7608736	0.76803344	12
3	13.09422077	12.5849625	12.9406511	0.35568864	12.5849625
4	9.819866178	12.9786264	13.2554774	0.27685099	13
5	7.855509279	13.2934997	14.0622147	0.76871502	13.3219281
6	6.546044604	13.5217687	14.4345582	0.91278945	13.5849625
7	5.610764895	13.7373074	14.2103024	0.47299496	13.8073549
8	4.909333659	13.950612	14.5915276	0.6409156	14
9	4.363792946	14.0461896	14.172074	0.1258844	14.169925

Comparing the minimum bits resolution achieved to those predicted in section 4.3 (calculated in the right hand column) indicates that the resolution limitation formula is

indeed correct, and that the simulation is behaving properly.

A subtle and initially disturbing item in the above chart is the average resolution obtained for N=1 of 11.36. As discussed earlier, average resolution for a phase lock condition, since each carrier zero crossing should have the same exact phase relationship with any single clock half cycle. Further investigation found the source of this error to be rounding in the spreadsheet. At two points during the simulation, the INT function swapped sides of the clock due to rounding, as the two are simulated to be exactly in phase.

By adjusting the initial carrier zero crossing time slightly, a more accurate picture of the average resolution minima is painted. Specifically, setting

$$zct_0 = \frac{1}{2^{20} f_c}$$

a revised table of minimum and average resolution achieved at the 9 N points is obtained, shown as table 5.3.

Table 5.3 DC Resolution at N Points - Part Two

N	Acceleration	Min Bits	Avg.	Difference	Predicted
1	0	11.0000055	11.0000055	2.4201E-08	11
2	19.64452956	11.9928511	12.0000378	0.00718671	12
3	13.09422077	12.5849789	12.584979	6.8486E-08	12.5849625
4	9.819866178	12.978648	13.0001383	0.0214903	13
5	7.855509279	13.2935266	13.3221345	0.0286079	13.3219281
6	6.546044604	13.5218002	13.585666	0.06386576	13.5849625
7	5.610764895	13.7373439	13.8081101	0.07076622	13.8073549
8	4.909333659	13.9506545	14.0004827	0.04982819	14
9	4.363792946	14.0462351	14.1721237	0.12588862	14.169925

The revised table clearly demonstrates that both average and minimum DC resolution obtained at N points are accurately predicted by the resolution limitation

equation. The slight deviations, increasingly apparent for higher values N, may be attributed to the finite length filter. Since a larger value of N results in a longer beat cycle, more filtering is required to obtain an average equally weighting each measurement within a beat cycleSpecifically, since N represents the number of carrier cycles per beat cycle, and F corresponds to the number of carrier cycles which are averaged, it is the relationship between N and F which determine how close DC resolution nears the optimum value defined by the resolution limitation equation. For an integer relationship between N and F (with F larger than N), we should see the minimum and average resolutions identical. For non integer relationships, the minimum resolution is slightly lower than the average, resulting from the partial beat cycle. This error should shrink as the ratio between F and N grows larger, as the partial beat cycle has less of an effect on the average.

As are result of this property, if rectangular filtering is used, and resolution limitation is dominated by a particular N frequency, it makes good sense to use a filter length which is a multiple of N (assuming that filtering is performed at the carrier zero crossing rate). This assumes, of course, that minimum resolution is the design metric of choice, since average resolution is not affected by partial beat cycles.

5.1.4 DC Behavior Near N Points

While the DC resolution attained at N points is limited by the resolution equation of section 4.3, DC phase ramp resolution immediately adjacent to N points is not. If infinitely long averages were taken, resolution at all points other than N and I points would be infinite. When a finite length filter is used, DC performance is degraded in the neighborhood of N and I points as demonstrated in the simulation results reported thus far.

Since frequency values in the close proximity of an N point frequency correspond to extremely large N point frequencies, the resolution at these frequencies should be extremely high. This notion is confirmed by the DC simulation results, with the notable exception that frequencies extremely close to an N point frequency have close to the same resolution as the N point frequency itself. The exact beat length at these frequencies is extremely high, but something very close to beating occurs over N cycles, with a slight phase shift each N cycles. This effect can be clearly seen by plotting DC resolution at frequencies in the immediate proximity of N points, with fines frequency steps. Consider, for example, figures 5.4 and 5.5, depicting DC resolution around frequencies where N=2 and N=3.

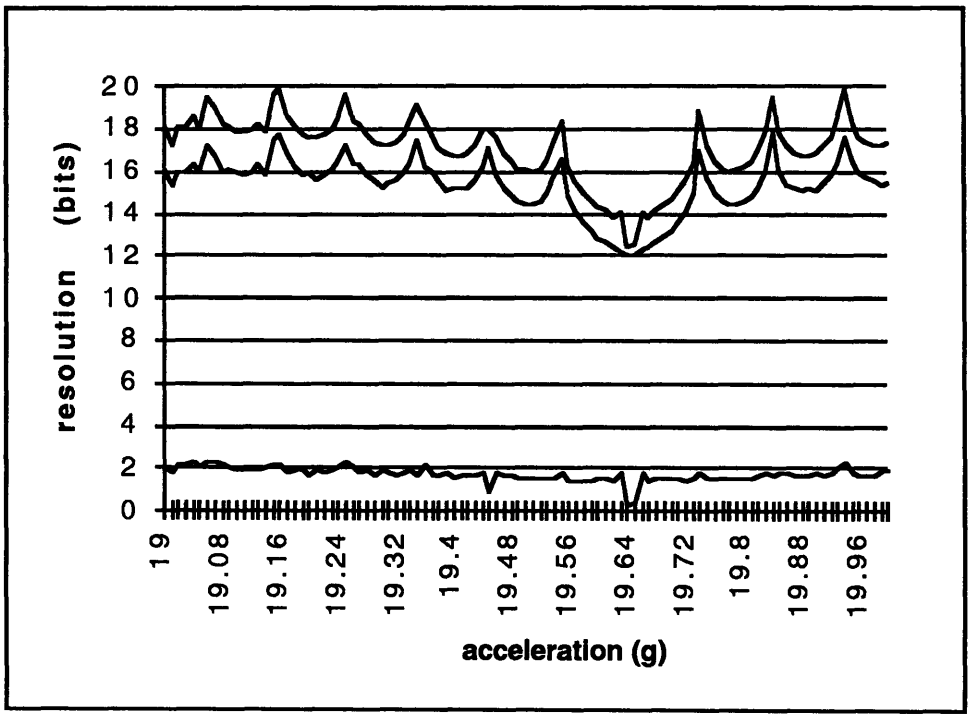

Figure 5.4 Resolution in the Vicinity of an N Point, N=2

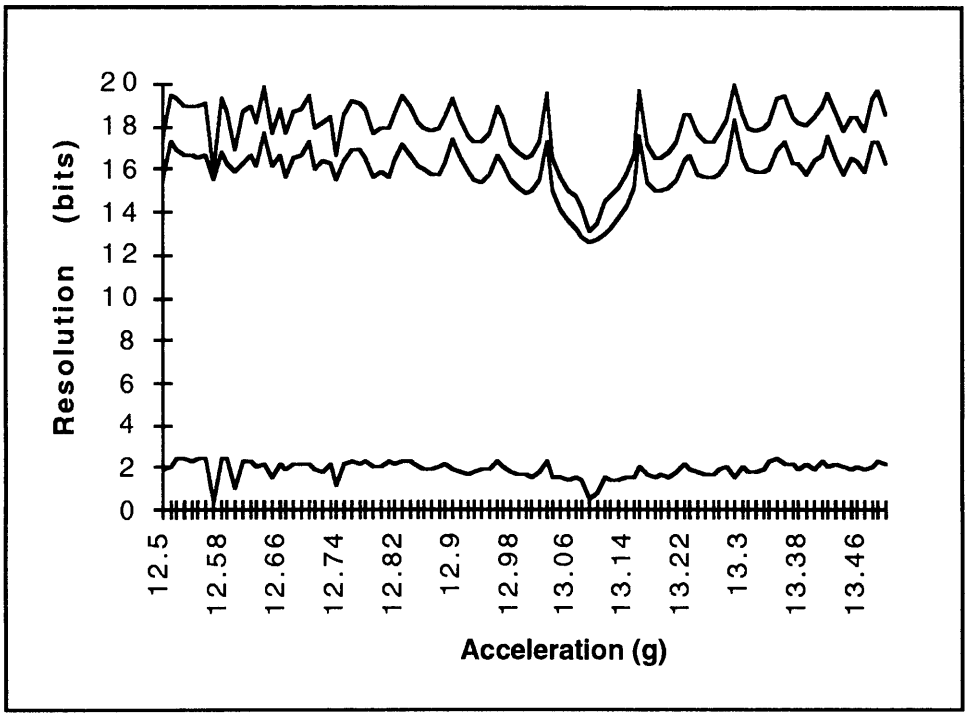

Figure 5.5 Resolution in the Vicinity of an N Point, N=3

Notice how the width of the resolution gap narrows for N=3. This region becomes exceedingly narrow for higher values of N, explaining the "masking" of N points in the previous simulation sets where frequency variations between simulations were much larger than here. This resolution gap region in the vicinity of an N point becomes exceedingly narrow with increasing values of N and F, and is bounded by resolution peaks on either side.

5.1.5 Average Resolution vs. Expected DC Accuracy Performance

Thus far, the average resolution measurements reported were taken as averages over many simulation cycles. For each simulation, corresponding to a point in one of the previous charts, a non-random initial phase relationship was chosen to demonstrate a "worst case scenario" so that resulting measurements correspond to the minimum values

which can be expected independent of phase angle.

A better indication of the performance which would probably be achieved for a certain set of clock and carrier frequencies may be an "expected DC accuracy" metric. Expected DC accuracy, considered as a function of acceleration/frequency/phase ramp slope once again, may be established by simulating over a range of initial conditions chosen from a uniform random distribution, or by averaging over a ramp of initial phase conditions. Using the latter approach, figure 5.6 was obtained, depicting resolution as a function of initial phase condition for N=2.

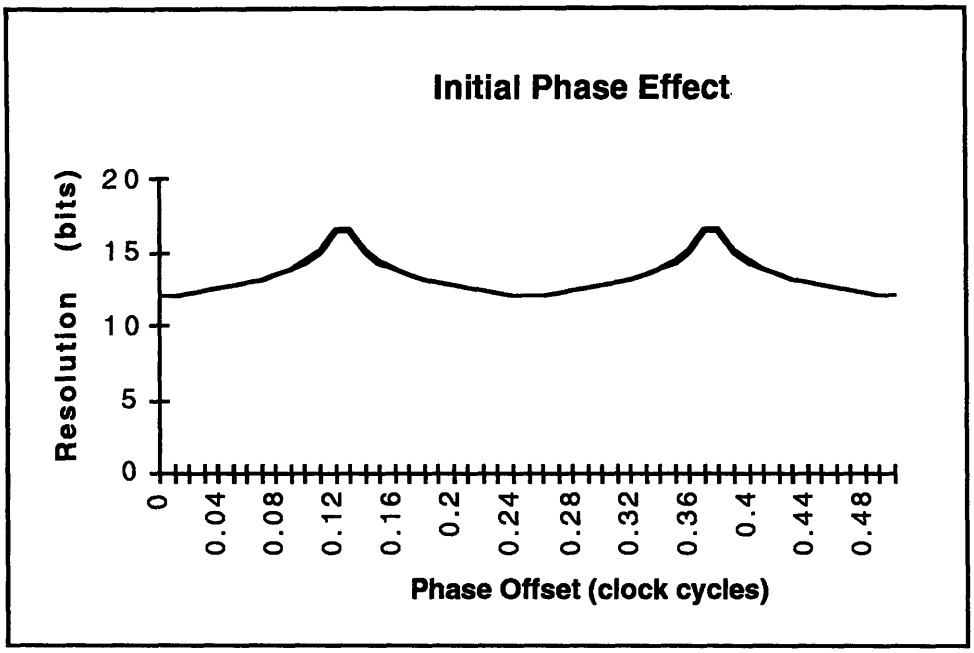

Figure 5.6 Phase Effect on Measurement Resolution at an N Point

Since N =2, two resolution variation patterns occur over half of a clock cycle instead of one. To compute a value for expected DC minimum resolution, an average over the phase offset variations results in 13.34 bits of expected minimum and average resolution, an improvement of 1.3 bits beyond the worst case values predicted by the resolution limitation equation and computed by the previous simulations.

Unfortunately, it can not be expected that this improvement in resolution will extend to all frequencies. This becomes readily apparent when a DC acceleration value is chosen arbitrarily away from an N point and a corresponding set of simulations are run. Such a case, with an acceleration of 12.9, is pictured in figure 5.7 showing both average and minimum resolution.

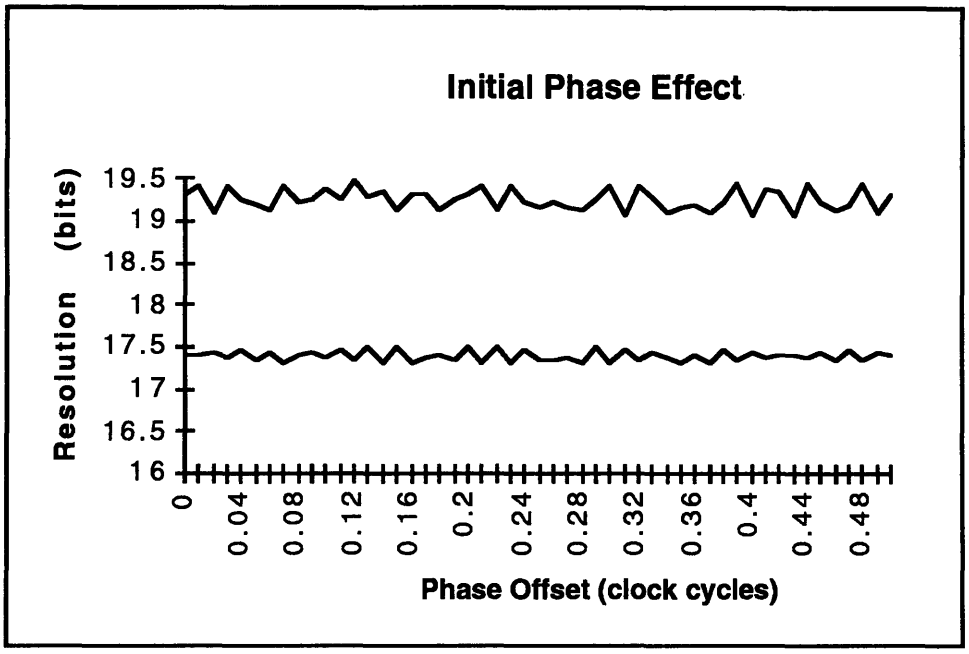

Figure 5.7 Measurement Resolution vs. Phase Away From N Points

Hence using an expected resolution metric improves resolution at N points by about a bit, and has no effect elsewhere in the spectrum.

5.1.6 Noise Effects on DC Resolution

All the simulation results presented thus far has been in the absence of noise. As indicated previously, the presence of noise ought to have a substantial affect on the DC resolution performance of the clock strobing readout.

A set of simulations were performed to determine the affects of noise on DC performance limitations. A substantial improvement in both average and minimum resolution was observed at N points when a small amount of noise was added. Consider the affect of noise at N=2, charted in figure 5.8 below.

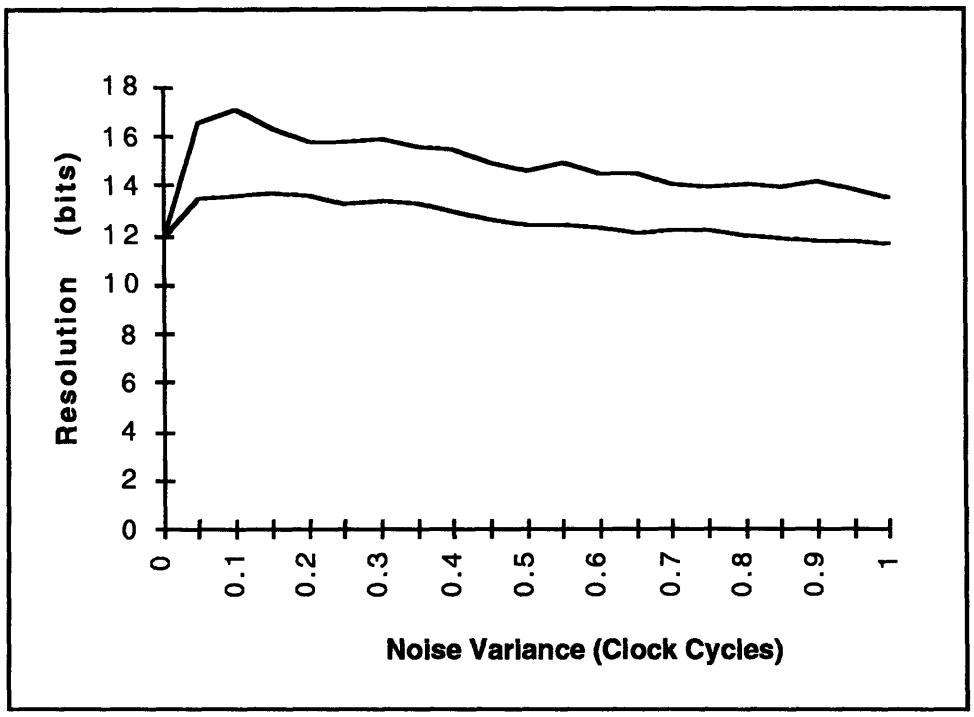

Figure 5.8 Effects of Noise on DC Measurement Resolution

For N=2, the optimal amount of phase noise variance is somewhere between .1 and .15 clock cycles, depending upon whether minimum or average resolution is of interest. This seems in line with the analysis offered previously, which indicated that for noise to improve measurement resolution, it would need to have a variance on the order of half the clock's resolution. In this case, the clock's resolution of 1/4 of a cycle (for N=2) corresponds to a noise variance of .125 clock cycles, resulting in an improvement of almost two bits of minimum resolution and 5 bits of average resolution near the noise detection threshold defined previously.

When the effects of noise and initial phase (considered in section 5.1.5) are considered previously, a better picture of the overall DC performance which might result under real life conditions. The following chart represents the minimum resolution achieved in a two dimensional simulation set where both the noise and initial condition parameters were varied in steps from one simulation to the next. The results are depicted for N=2 in figures 5.9 below and 5.10 below. The perspective in figure 5.10 is adjusted to show the affects of increasing noise amplitude.

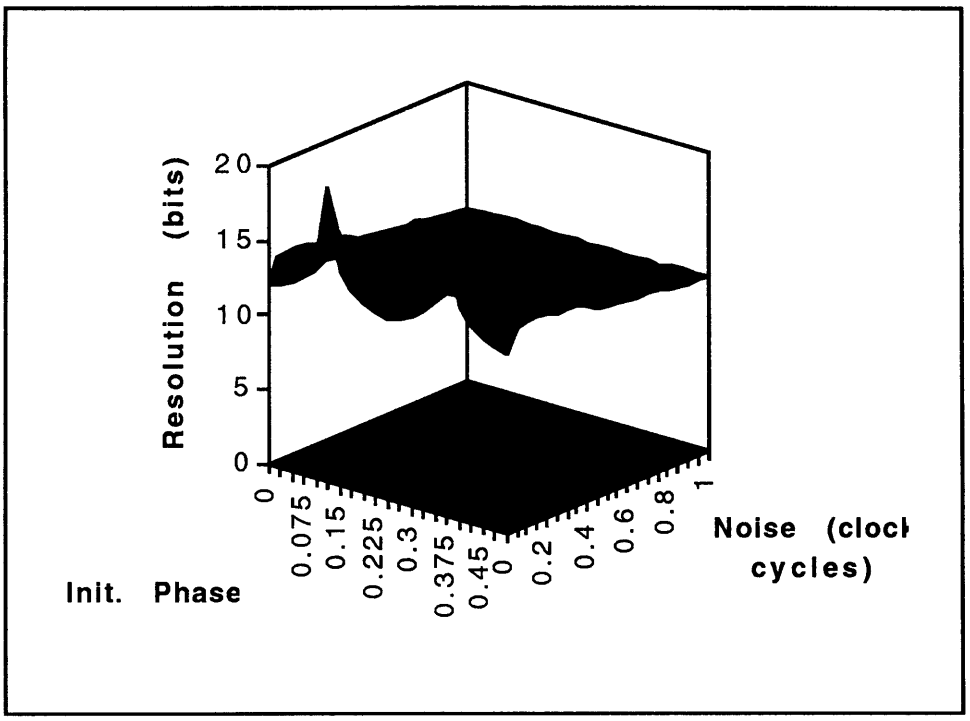

Figure 5.9 Noise and Initial Phase Effects on DC Min. Resolution at N=2

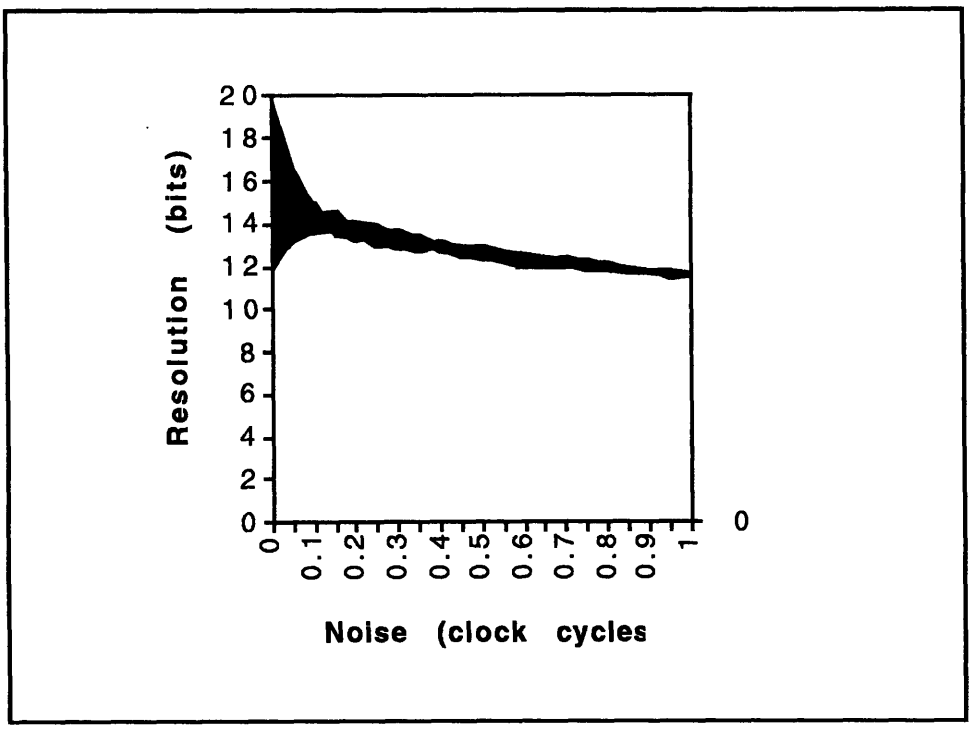

Figure 5.10 Effect of Noise on DC Min. Resolution at N=2

In two dimensions, the effect of noise can be seen as raising the minimum resolution (and average resolution, not depicted) to the expected values. Hence the sides view above clearly shows a narrowing in initial phase deviations as the noise detection threshold is reached. Whether or not resolution improvements substantially beyond expected value are possible was not established.

5.2 Dynamic Performance Characteristics

Several sets of simulations were performed to establish performance limitations on the clock strobing readout architecture in a dynamic environment. Parameters varied over the simulations included signal characteristics, noise, clock to nominal frequency ratio, and filtering characteristics.

The point of this work was not to characterize the readout as a linear system. First, it is not a linear system. The time quantization function of the clock is clearly nonlinear. Second, any dynamics induced by the readout architecture are not a function of a clock strobing readout inherently but of the particular filtering implementation used. The purpose of dynamic environment simulation was to determine what extent resolution inherently improves or declines from the DC limits in a changing environment.

5.2.1 Step Response Performance

A set of simulations were performed to establish the performance limitations of the clock strobing readout when a step of acceleration (frequency) occurs. Initial and final accelerations were set, and the step set to occur midway through the simulation. Consider the step response charted in figure 5.11 below.

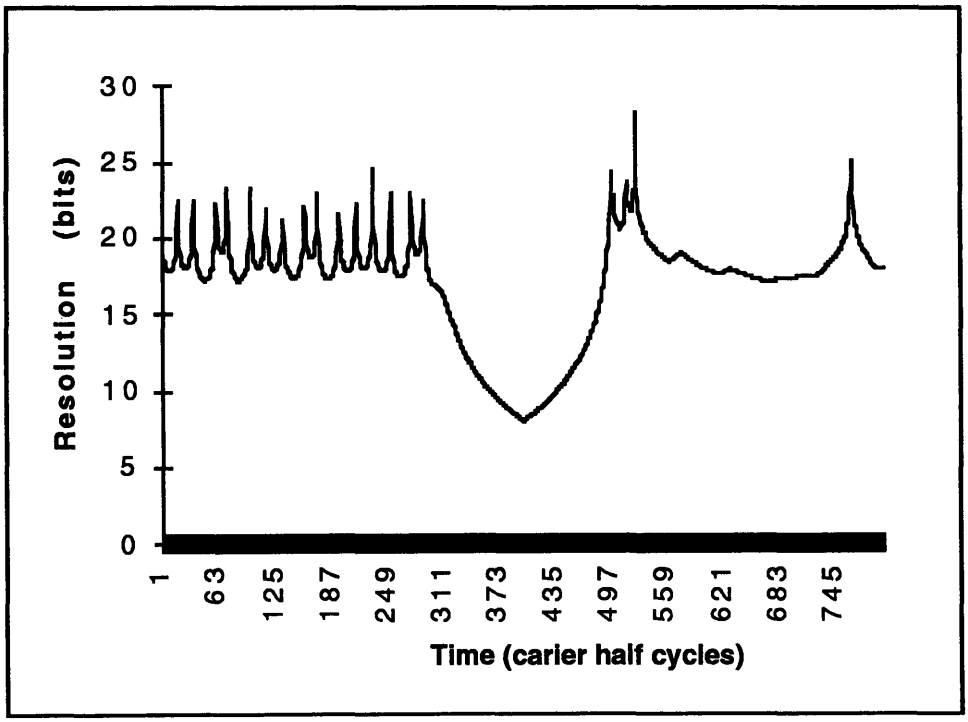

Figure 5.11 Acceleration Step Response

In this case, a triangular 201 point filter was used. Initial and final acceleration values of 1 and 10 g were set, along with the same clock and carrier frequencies used in the previous examples. No noise was present.

As would be expected, when the step is beyond the filter's reach, resolution matches DC resolution. Resolution changes occur within 100 cycles of the transition in either direction. Resolution bottoms out at 8 bits, right where the acceleration step occurs. Notice the resolution changes after the step occurs, due to the new acceleration value and corresponding clock to carrier frequency ration

It should be expected that a lowpass filter will respond poorly to a step input, as indicated by the performance above. In the case where a bandlimmited phase modulated signal is considered, steps will never occur, and the step response of the DCSR is of no real interest, since error is a nonlinear function of the input signal, and no general model of readout behavior can be extrapolated from the step response..

If frequent non-band-limited steps in input frequency is expected, then either no filtering should be used, to prevent average resolution from dropping below unfiltered values, or a non-linear adaptive filter could implemented which stops averaging in a region when actual measurements correspond to apparently real phase changes beyond what can be accommodated by the filter.

5.2.2 Response to Sinusoidal Acceleration Profiles

Resolution response of the DCSR system to sinusoidal acceleration profiles was simulated for a range of frequencies within selected frequency bands. As discussed previously, since the readout is non-linear, the applicability of these results to what can be expected when multiple frequencies are present is limited. Nonetheless, response to

sinusoidal signals was seen as an important first step toward understanding what resolution performance may be achieved by the readout in a changing environment.

Consider the resolution of a 50 Hz sinusoidal acceleration profile with an amplitude of 1 g., using a 201 point triangular filter, as depicted in table 5.4 for multiple values of N.

Table 5.4 Resolution Performance for Sinusoidal Modulation - Part 1

N	Min. Res.	Avg. Res.	Unfilt. Avg..
1	10.66898463	11.92371291	12.03
2	11.00722485	12.10662058	12.66
3	11.09686439	12.14851906	12.37
4	11.1474756	12.16738165	12.49
5	11.15136676	12.16348854	12.44
6	11.11989946	12.15564107	12.46
7	11.1754275	12.15907333	12.44
8	11.19171253	12.16816813	12.39
9	11.16028179	12.15296641	12.46
10	11.14393497	12.16057752	12.47

In this particular set of simulations, the effects of excessive filtering are readily apparent. Because frequency variations resulting in greater than a the clock's resolution occur over 201 carrier cycles, the filter actually degrades resolution performance. When the filter length is shortened to 21 points (rectangular), improved resolution is visible, as depicted in table 5.5.

Table 5.5 Resolution Performance for Sinusoidal Modulation - Part 2

N	Min Res.	Avg. Res	Unfiltered
1	11.12411089	12.6773885	12.0329
2	12.17762166	14.95674139	12.6597
3	12.58401132	15.32251876	12.3672
4	13.36623529	15.74586834	12.4932
5	13.07627362	15.84663721	12.4375
6	13.06863582	15.98112346	12.4617
7	13.1873986	15.77051067	12.4397
8	13.66614917	16.00264717	12.3885
9	13.24902517	15.95089613	12.4628
10	13.21876201	15.57040991	12.4692

For increasingly large values of N, however, resolution does not improve beyond 13 or so bits. Overall, the resolution achieved here is not dramatically lower than what was seen for DC signals when the 21 point filter was used.

To determine the bounds in frequency and amplitude of an acceleration where the DC resolution equations still approximately hold, a group of two dimensional simulation sets were performed. Several of the resulting plots are included at the end of Appendix A. Consider the minimum and average results shown in figures 5.12 and 5.13, respectively, charted for N=4.

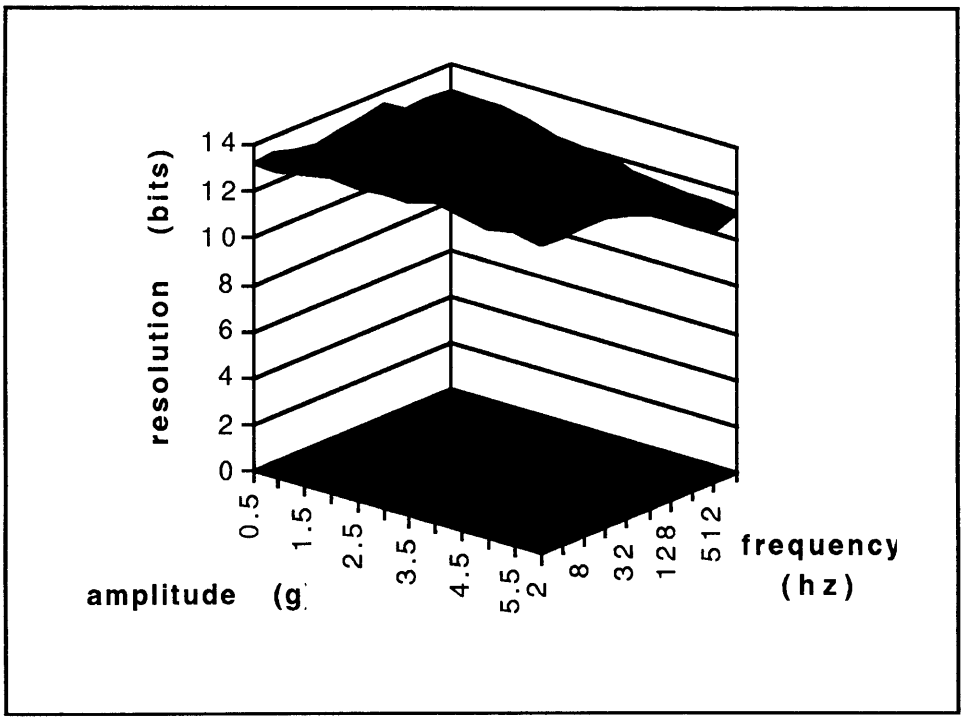

Figure 5.12 Minimum Dynamic Resolution Near N=4

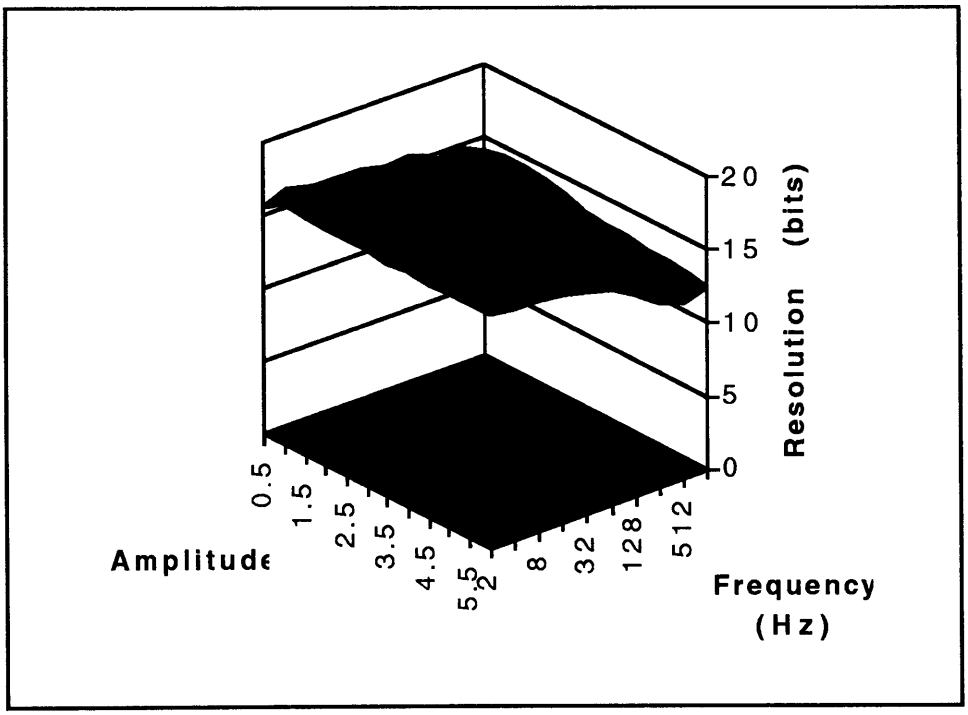

Figure 5.13 Average Dynamic Resolution Near N=4

Both average and minimum resolution performance stays fairly flat at the N point DC resolution of 13 across acceleration and frequency over a fairly wide range, ending at a modulation frequency of about 64 Hz.

The same simulation set was also run around an I point, yielding the distinctively different results shown in figure 5.14 and 5.15 below, for minimum and average resolution respectively..

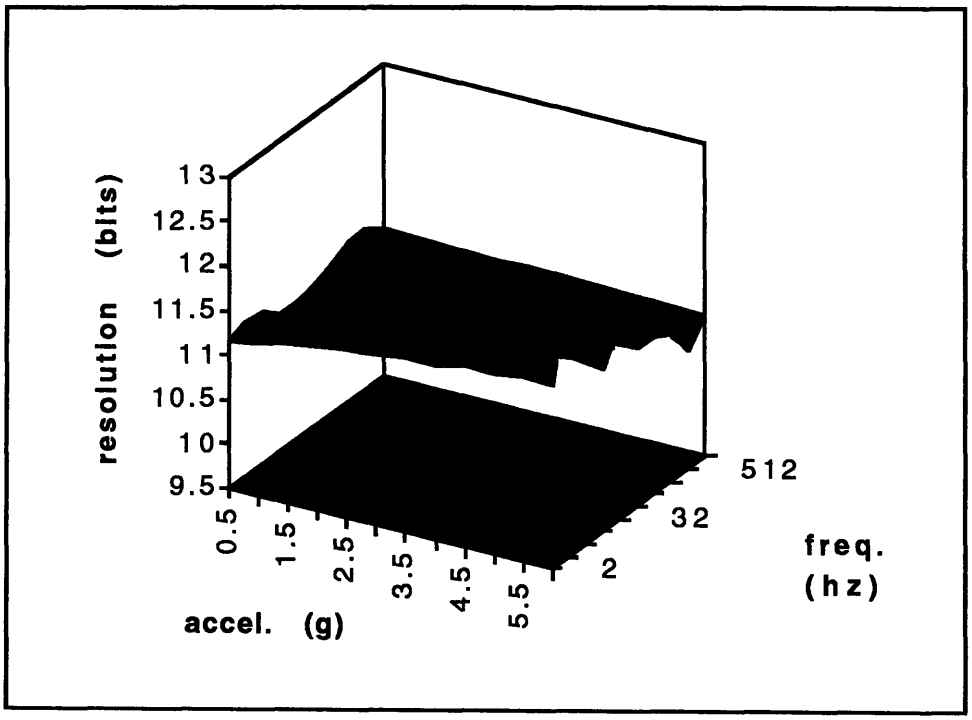

Figure 5.14 Minimum Dynamic Resolution Near an I Point

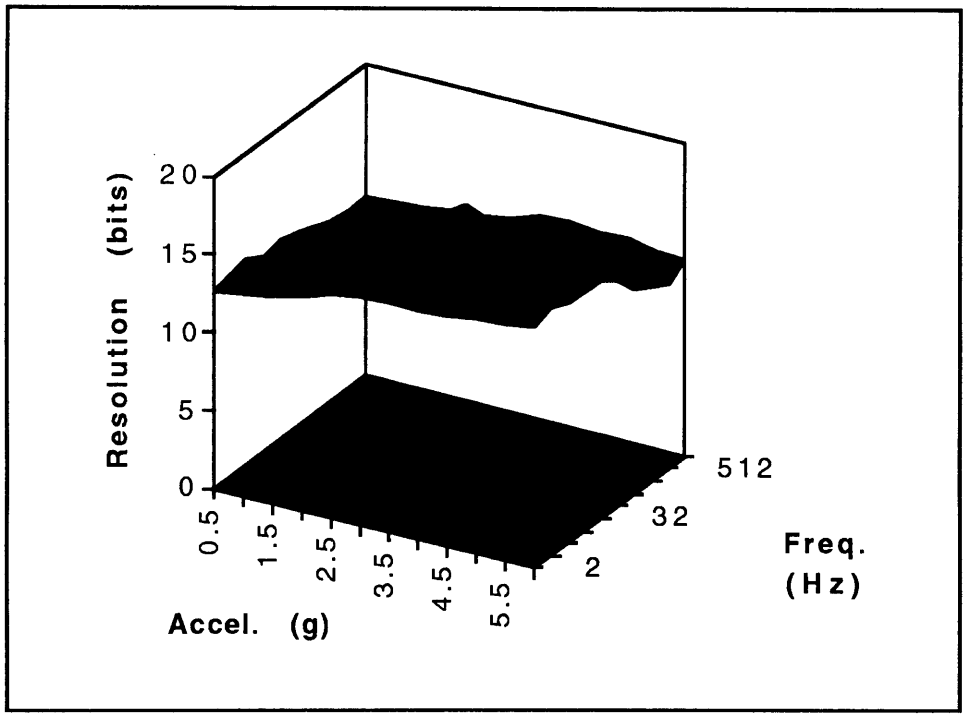

Figure 5.15 Average Dynamic Resolution Near an I Point

The I point simulation demonstrates how modulation can actually improve resolution beyond that suggested by the resolution limitation equation for sufficiently low modulation frequencies.

After running many sets of sinusoidal acceleration modulation simulation sets, no pattern or formula became clear for dynamic resolution, though several notions were indeed confirmed. First, for sufficiently slow variations in acceleration and phase, DC resolution limitation equations provide an accurate prediction of performance. Second, a "safe" range of modulations seems to exist, though not well defined. With a filter length of 1/10 to 1/20 of a modulation period, and values of N on the order of the square root of the filter length, dynamic resolution achieved the values predicted by the DC resolution limitation equations. Third, for low values of N or at I point frequencies, modulation may result in resolution improvements, for sufficiently low modulation frequencies.

Response to Band-limited Signals

Several Simulations were performed using MATLAB generated bandlimmited signals. The results of these simulations did not contrast with the results of sinusoidal experiments, though a much smaller range of band limits and amplitudes was tested. Since the spreadsheet simulation was not easily modified to generate bandlimited signals internally, performing a large number of bandlimited simulations proved impractical. Enough points were tested to strongly indicate that the sinusoidal results are typical of time varying input response of the Digital Clock Strobing Readout, and that bandlimited performance may be conservatively approximated by sinusoidal performance at the band limit.

5.3 Conclusions

Simulations accurately confirmed the resolution and accuracy performance of the Digital Clock Strobing Readout. DC resolution at much higher than the clock's resolution may be obtained for an appropriate clock to carrier frequency ratio. Overall DC resolution over a broad spectrum of frequencies improves an average of 2.5 bits per decade, but improves as high as 3.5 bits per decade at some frequencies, and improves none whatsoever for N and I point frequencies. Time varying signals will result in corresponding DC resolution values for sufficiently slow time variations. Given a frequency deviation ratio which places the carrier in a range of frequencies with a dominant (minimum) value of N, and a bandlimited modulation, if filtering occurs over a small fraction of a band limit frequency modulation cycle, then a factor of N improvement in resolution can occur for N on the order of the square root of the filter length, potentially higher.

For frequency relationships which are plagued by I or N points within the range of a signal variations that prevent enough resolution improvement to take place, several options are available. In such a case, increasing the clock's resolution by synchronizing additional bits may be appropriate. A second approach would be to synchronize additional bits to the carrier itself. A third approach would be to use a third frequency reference (2nd clock) which would not necessarily need to be any higher in frequency than either of the two original waveforms (especially if the resolution limit is induced by a particularly bad combination of the original frequencies, less so if the limitation is due to wide or rapid modulations of the carrier). Any of these approaches demands a thoughtful choice in terms of what frequencies should be generated from the existing ones to optimally make measurements. A specific example of such a choice being made is offered in section 6.2. Some related ideas for improving effective resolution are discussed briefly in Appendix B.

Chapter 6: DCSR Analysis Applied to Other Phase Readout Approaches

An important aspect of the simulation and analysis performed for the Digital Clock Strobing Readout is that the results are applicable to wide range of readout architectures. In particular, all of the DCSR analysis applies to digital phase locked loops with zero crossing phase detectors. With respect to analog digitizing phase locked loops, virtually all of the DCSR results transform into an equivalent set of relations to determine the resolution limitations for this architecture. With respect to the Delta Sigma Phase Locked Loop, resolution is limited by the same equations if synchronization is performed prior error detection, or if phase error is quantized by the clock.

6.1 Digital Phase Locked Loop Resolution Limitations

The resolution limitation equations apply not only to the Digital Clock Strobing Readout but to the entire class of digital phase readouts that rely that effectively "round" carrier zero crossings clock resolution increments. The class of digital phase locked loops which demodulate in this fashion is referred to by Lindsey and Chie as Flip-Flop Digital Phase Locked Loops, because flip flops are often implemented as phase detectors in these implementations. These DPLLs all suffer from the same fundamental resolution

constraints as the DCSR. The origin of this error may be attributed to the inherent loss of information which occurs when a phase lock condition occurs with a finite error, at an N point or I Point frequency. Until the error grows larger than the effective resolution of the clock at this frequency, no measurement improvement can occur No amount of filtering can recover information which no longer exists.

In fact, as became clear when zero crossing time processing was considered in section 4.4, any "all digital" phase locked loop's performance may be realized from the zero crossing times of the DCSR, including optimal filtering, without regard to feedback stability issues which might arise from high order filters in a close looped system such as a digital phase locked loop. This includes the Delta Sigma phase locked loop as well, when implemented with a digital phase detector. In this case, no matter how optimally the quantization noise is shaped, the information which is being modulated has uncertainty which cannot be eliminated.

With respect to other types of digital phase locked loops, the fact is they are not purely digital. They require the use of an analog phase detector, though this may or may not be a concern.. When an analog phase detector is used, the DC resolution limitation equations no longer apply for any value of I or N. Instead, the frequency and phase resolution limitations are determined by the NCO, loop filter, and interpolation implementation. The effects of fedback phase quantization noise does not affect DC resolution so long as the loop keeps an average of zero phase error. Quantization error induced by finite phase or frequency steps will need to be filtered out, but no DC component will exist, resulting in resolution potentially much higher than the clock's resolution for sufficiently slow varying signals of any frequency. Several implementations of this sort are discussed in Appendix C.

6.2 Analog Phase Locked Loop Resolution Limitations

Interestingly enough, while the Digital Clock Strobing Readout could be seen as a design alternative to the use of analog phase locked loops, the two approaches are surprisingly identical in nature, performance, and underlying principle. As discussed in section 4.1, when only two time references are available, it doesn't matter which signal strobes the other. In fact, the Digital Clock Strobing Readout Simulation can be viewed as a Phase Locked Loop Strobing simulation, if the clock is considered the instrument, and the instrument the clock. This analysis holds perfectly except to the extent that the filter bandwidth is no longer modulated by the carrier modulation. This distinction is small and has been neglected thus far, so I will continue to do so.

The resolution limitations of the analog digitizing phase locked loop become clear when certain equivalencies are made, enumerated as follows:

1) The rate at which the digitized phase is strobed by the system clock corresponds to half carrier frequency in the previous analysis. The factor of two corresponds to the fact that a carrier cycle comprised of two half carrier cycles, distinguished by up and down zero crossings. Whatever rate strobing occurs corresponds to the rate carrier zero crossings occur in the clock strobing readout.

2) Noise is no longer a factor which affects strobed measurements (to the extent that noise has been filtered by the phase locked loop).

3) The VCO now corresponds to the clock. The highest frequency bit of the clock corresponds to the VCO frequency itself, and each of the others corresponds to bits in the feedback counter.

4) The frequency relationship of interest to determine resolution limitations is now half the VCO to strobe frequency ratio. N and I are defined as before in terms of these frequencies. To make the translation complete, we can define a "new" resolution limitation equation from the frequency relationship

$$I\frac{f_s}{2} = Nf_v$$

where

$$f_v \equiv vco\,frequency$$

and

$$f_s \equiv strobing\,frequency.$$

The resolution limitation equation

$$b_c = \log_2(I) + 1$$

holds as before, but in this case b_c corresponds to bits resolution with respect to half the strobing frequency. To convert bits with respect to the strobing frequency to bits with respect to the incident carrier at the loop's input (f_n, nominally), we can use the relationship

$$b = b_c + \log_2\frac{f_s}{2f_n} = b_c + \log_2\frac{M \times f_s}{2f_v}$$

where

$$M \equiv loop\ feedback\ counter\ length.$$

In light of these equivalencies, the resolution performance of the analog digitizing phase locked loop can be determined by analysis or by simulation as before. The first relevant question to ask is how much could the resolution of an analog digitizing phase locked loop be improved without the addition of an additional VCO, or a higher frequency VCO?

Consider the SFIR example once again, using an analog digitizing phase locked loop readout with a VCO that oscillates at a frequency 1024 times the 6.4 kHz carrier ($M = 1024$). The VCO is strobed at a rate of 6.4 kHz ($f_s = 6.4\ kHz$). Since half the strobing frequency is an I Point frequency of the VCO frequency ($I = 2048$), the resolution limitation equation indicates that no amount of filtering will result in measurements with DC resolution better than

$$b = \log_2 \frac{I \times M \times f_s}{2f_v} + 1 = \log_2 \frac{2048 \times 1024 \times 6.4kHz}{2 \times 1024 \times 6.4kHz} + 1 = 11$$

bits. It seems appropriate, however, to consider strobing the loop at a rate higher than 6.4 kHz to get more measurements, considering that clock bits are available at much higher frequencies. Consider the effect of strobing at 25.6 kHz. In this case, every parameter is the same as previously with the exceptions $I = 512$ and $f_s = 25.6\ kHz$.. The resolution, unfortunately, does not change since the product $I \times f_s$ does not change, producing an equivalent 11 bits of phase resolution with respect to a nominal carrier cycle. For each multiple of two the strobing climbs in frequency, the $I \times f_s$ product will remain the same.

Suppose, however, that the phase digitizing feedback counter in the loop counted down by a factor of 1025 instead of by 1024 ($M = 1025$). Once again, if strobing occurs at the rate of 6.4 kHz, a phase resolution of almost exactly 11 bits . This time, however, when the strobing rate increasing by a factor of four, the frequency relationship considered for determining resolution limitation is defined by

$$(1025)\frac{24.8\ kHz}{2} = 2 \times (1025 \times 6.4\ kHz)$$

corresponding to values of $I = 1025$ and $f_s = 25.8kHz$, resulting in 12 bits of phase resolution, a factor of two improvement. This improvement occurs because the second factor of two increase in strobing rate results in an N of two, rather than dividing I by two as in the previous example (where I stayed at 512). In fact, each additional factor of two increase in strobing rate will result in a factor of two improvement of resolution[1].

When considering the filtering burden to obtain the maximum resolution for large values of N, it is comforting to realize that filtering over N cycles corresponds to N cycles of the strobing waveform who's frequency is proportional to N. Hence, the amount of time which filtering must take place is never more than one half of one 6.4 kHz cycle, in this particular case. And in general, raising the sampling rate by a integer multiple to obtain higher resolution measurements does not result in a longer filtering time required.

To improve resolution by increasing strobing rate, the necessary criterion, as demonstrate above, is that I not be divisible by N. In the SFIR example, with the 1024 point counter, doubling the strobing rate (which is all that can be done when all of the clock bits are related by frequency powers of two) provides no advantage because I=1024 is a power of two.

In general, all of the Digital Clock Strobing Readout analysis is a useful

[1]The Hypha Report (Cox et al) demonstrates how to achieve this effective increase in readout resolution, actually using I=1023. The effects are explained in terms of a Vernier scale.

methodology for analyzing and predicting the performance of the analog digitizing phase locked loop, with noise affects being the primary difference.

6.3 Conclusions

Since the DCSR analysis applies in general to measuring the phase relationship between two signals when both signals are available in discrete evenly spaced phase increments (though they may vary over time), the results may be used to derive the inherent resolution performance limits for any combination of frequencies. Hence the resolution limits for both digital phase locked loops and analog digitizing phase locked loops, as well as any direct clock strobing approach, may be determined from the results of Chapters 4 and 5, or by using the DCSR simulation with an appropriately chosen filter.

Chapter 7: Concluding Remarks

While the study of measurement and digitization techniques for narrow band phase modulated signals is far from complete, the analysis and simulations performed and reported herein clearly demonstrate several important principles with design implications. In particular, the following observations should be duly noted.

1) Resolution substantially beyond the highest VCO or clock bit frequency in a phase readout a digitization scheme is possible and in general should be expected.

2) For some signals, resolution beyond the highest clock or VCO bit cannot be obtained, due to phase lock conditions.

3) The maximum achievable DC resolution of any zero crossing detector based Digital Phase Locked Loop, or analog digitizing phase locked loop, may be determined from the same fundamental resolution limitation equations, with average and expected performance generally a few bits higher.

4) While a definitive formula for phase resolution limitations was not found for dynamic phase modulated signals, a rough set of conditions which justifies the use of DC

resolution limitation equations in a time varying environment may be used.

5) For slow modulations at frequencies which do not guarantee enough resolution in a particular application, the controlled addition of just the right amount of zero mean phase noise can substantially improve performance in any of the all digital techniques, or the analog digitizing phase locked loop. This approach will be particularly successful if the noise has a high frequency spectrum.

6) If all digital techniques using existing frequencies do not provide the adequate resolution for a particular set of available frequencies, then additional frequencies may be used, higher frequencies generated, or hybrid techniques implemented to gain the desired resolution.

7) The Digital Clock Strobing Readout method may in certain situations offer improved performance, more practical implementation, and ease of use than it's digitizing phase locked loop counterpart. If a high frequency version of the carrier exists, than most likely all of the bits of both frequencies should be used to obtain optimal performance. If either a signal or the clock needs to be beat up to a high frequency, that in certain respects it makes more sense to synchronize higher frequencies to the clock. This is particularly true if the digital signal processing power is readily available, and the use of analog components is to be minimized.

Appendix A: Clock Strobing Simulation

This simulation investigates the potential resolution and accuracy of the digital clock strobing readout architecture. The spreadsheet is designed to model phase readout performance during some number of carrier cycles. The input to the simulation is a column of acceleration values. Since each row represents one half carrier cycle. These half cycles are stretched or squeezed in time with respect to nominal values under non-zero accelerations. Hence a row represents a variable amount of time dependent upon the exact acceleration during that cycle.

Input to the Simulation

To simulate the response of the clock strobing readout, several parameters are used. Namely, a digital clock frequency, a nominal carrier frequency, and an instrument scale factor are provided. These parameters are represented as follows:

$$f_c \equiv \textit{digital clock frequency (hz)}$$

$$f_n \equiv \textit{nominal carrier frequency (hz)}$$

$$sf \equiv instrument\ scale\ factor\ (radians / sec / g)$$

so that

$$\phi_i(t) = f_n t + \theta(t) \quad (cycles)$$

where

$$\phi_i(t) \equiv instantaneous\ signal\ phase$$

and

$$\theta(t) \equiv \ relative\ phase\ angle$$

calculated as

$$\theta(t) = 2\pi \int (sf \times a) dt$$

Simulation Time Steps

Each row in the spreadsheet simulation corresponds to a modulated carrier zero crossing. Hence the time elapsed between each row of the simulation is approximately one half of one carrier cycle:

$$\Delta t \approx \frac{1}{2 f_n}$$

where Δt is the time elapsed between any two rows. The exact time between rows is the amount of time between zero crossings of the modulated carrier, rather than the unmodulated carrier and therefore varies by some amount determined by the acceleration during the interval and the instrument scale factor:

$$\Delta t_n = \frac{1}{2(f_n + \frac{sf \times a_n}{2\pi})}$$

where

$$\Delta t_n \equiv \textit{time elapsed between zero crossings } n\text{-}1 \textit{ and } n$$

and

$$a_n \equiv \textit{average acceleration over } n\textit{'th half modulated carrier cycle}$$

In a typical case, such as in the model for the SFIR output signal, the variable component of the time interval varies by a small amount relative to the fixed component, so that the second term in the denominator of (Δt equation above) can be approximated as zero for certain calculations.

In general, for each variable considered during the n'th time interval, the subscript n is used to denote the proper interval, with the understanding that the sequence of variables over n denotes a stream of values taken at non uniform intervals of width Δt_n.

Input of Acceleration Profile

Acceleration values for each carrier half cycle may be either input directly into the spreadsheet or generated to simulate an approximately sinusoidal acceleration input. Should a sinusoidal acceleration input be simulated, an input frequency and amplitude must be provided as well. These parameters are represented by the variables

$$f_a \equiv \textit{frequency of sinusoidal acceleration (hz)}$$

$$A \equiv \textit{amplitute of sinusoidal acceleration (g)}$$

such that

$$a(t) = A\cos(2\pi f_a t)$$

where

$$a \equiv \textit{input acceleration (g)}.$$

Generation of Zero Crossing Times

The next column of the simulation calculates a vector of zero crossing times as a function of the input acceleration vector, the nominal carrier frequency, and the instrument scale factor. As previously mentioned, each row in the spreadsheet represents the state of the digital clock strobing readout following the detection of a single carrier zero-crossing time. The time which this zero crossing occurs is calculated as the sum of the previous

zero crossing time, plus a calculated period value determined by the average acceleration experienced by the sensor during the period.

$$zct_n = zct_{n-1} + \frac{1}{2} p_n$$

where

$$zct_n \equiv \textit{time of n'th carrier zero crossing (seconds)}$$

and

$$p_n \equiv \textit{period of nth cycle}$$

calculated as

$$p_n = \frac{1}{f_o + (sf \times a)}.$$

Calculation of Measured Zero Crossing Times

From these "actual" zero crossing times, a vector of "measured" zero crossing times is calculated, corresponding to the quantized measurement which is obtained by strobing the system clock in the digital clock strobing readout. The simulation calculates the zero crossing times measured by the readout from the actual zero crossing times and clock frequency. The simulation begins at t=0, and for any corresponding time in the future, a clock reading may be calculated as the actual time, quantized by clock resolution. Recalling that:

$$f_c = \textit{highest frequency of digital clock}$$

and defining

$$mzct_n \equiv \textit{measured zero crossing time of n'th zero crossing}$$

the measured zero crossing time is calculated as:

$$mzct_n = \text{int}(zct_n \times 2f_c) / 2f_c + (1/4f_c)$$

where the "Int" function takes the integer component of the n'th zero crossing time multiplied by twice the clock frequency to determine the number of half clock cycles which have passed since the start of simulation. The number of half clock cycles, divided by the time per half cycle, corresponds to the time indicated by the clock at the time of the n'th zero-crossing. This calculation implicitly sets the clock and carrier phases in perfect lock at time zero. A quarter clock period of time is added to each clock reading so that each measurement is accurate within plus or minus one quarter of a clock period of the actual zero crossing time rather than within minus one half of a clock period.

Addition of Phase Noise

To simulate the effects of phase noise on the clock strobing readout, a noise term was added to the measured zero crossing time formula. While in reality, noise would be reflected in actual zero crossing times, I wanted to preserved the integrity of the "actual" zero crossing times to reflect "ideal" zero crossing times for error generation. Adding noise to the measured zero crossing times within the quantized term equivalently simulates the effect of zero crossing times in the presence of noise, and maintains the

integrity of the actual zero crossing times for error measurement. Hence measured zero crossing times were calculated in these simulations as

$$mzct_n = \text{int}\left[\left(zct_n + \frac{K_{noise} \times R_n}{f_c}\right) \times 2f_c\right] / 2f_c + (1/4f_c)$$

where the third term is the noise term, comprised of

$$K_{noise} \equiv \textit{noise amplitute coefficient},$$

a simulation variable input, and

$$R_n \equiv \textit{normally distribute random variable,}$$

with a variance of 1 and a mean of 0 (generated in MATLAB and imported into a column of the spreadsheet). This term is divided by the clock frequency so a unity value of the noise amplitude coefficient. represents 1 clock cycle of noise variance.

Generation of Relative Phase Readings

For each zero crossing time measurement, a corresponding *relative phase measurement* is generated corresponding to the cumulative phase difference between the modulated carrier and unmodulated carrier of nominal value. Hence, the simulated readout calculates how much the measured signal has advanced or retarded in phase due to non-zero accelerations. The measured relative phase, defined as

$$\phi_n \equiv \textit{measured relative phase}$$

is calculated by:

$$\phi_n = (mzct_n - nzct_n) \times f_n$$

where

$$nzct_n \equiv \textit{calculated time of n'th nominal carrier zero crossing (sec)}$$

and

$$nzct_n = \frac{n}{2 \times f_n}.$$

Phase Error Calculation

For each simulated relative phase measurement, an error is calculated and converted to number of bits phase measurement resolution. Simulated phase measurements are compared with actual relative phase values, calculated as:

$$\theta_n = (zct_n - nzct_n) \times f_n$$

to generate error measurements, denoted

$$e_n \equiv \textit{measurment error (cycles)}$$

where

$$e_n = (\phi_n - \theta_n).$$

The number of bits measurement resolution, denoted

$$b_n \equiv \textit{equivalent bits measurment resolution}$$

are calculated as

$$b_n = \log_2\left(\frac{1}{e_n}\right) - 1$$

Where the subtraction on one bit is to account for the fact that with n bits of phase resolution, the maximum error ever seen would have a magnitude less than or equal to $1/(2^{n+1})$ with a range (positive and negative) of $1/2^n$. Hence 4 bits of resolution corresponds to phase measurement resolution of 1/16 of a carrier cycle, which corresponds to a maximum error of +/-1/32 of a cycle.

Filtering of Phase Measurements

Phase measurements are digitally filtered by processing the column of phase measurements. Defining a formula for the filtered value in a single row in terms of either the terms of the impulse response (for an FIR type filter) or a recursive formula (for an IIR filter) results in a fully filtered column of values when the value is copied into the rest of the rows of the spreadsheet simulation.

Generation of Sinusoidal Acceleration Profiles

If a single frequency input is being simulated, then the accelerations are calculated as

$$a_n = A\sin(2\pi {f_s}/{2f_n}\, n)$$

The resulting acceleration profile is approximately sinusoidal. An exact simulation of a sinusoidal acceleration profile requires that

$$a_s \equiv \frac{\int_{zct_{n-1}}^{zct_n} A\sin(2\pi f_s t)}{(zct_n - zct_{n-1})}$$

which is dependent upon the carrier zero crossing. Since carrier zero crossing times are calculated from acceleration values, an exact sinusoidal simulation proves more difficult.

The error in the simulated non-zero amplitude sinusoidal accelerations can be attributed to two independent factors. When simulation time step intervals vary for non-zero accelerations, the resulting acceleration profile "dithers" in frequency by a small amount depending upon the amplitude of the real, simulated accelerations. A second source of error is introduced because the simulated acceleration corresponds to a value at the beginning of the carrier half cycle, rather than an average over the completed cycle. None the less, given the conditions:

$$f_s << f_n$$

and

$$\frac{sf \times A}{2\pi} << f_n$$

it can be concluded that the spectrum of the simulated acceleration profile is a close approximation to a physical sinusoidal spectrum.

Appendix B: Hybrid Clock Strobing Approaches

A number of ways to improve the resolution of the Digital Clock Strobing readout using hybrid digital/analog techniques have been suggested These hybrid approaches, in one way or another, seek to improve the clock's resolution by the introduction of analog circuitry. Clock resolution may be improved through a number of techniques, several of which will be briefly discussed.

One technique, implied in Chapter 4 as an alternative to using a digitizing analog phase locked loop when not enough clock bits are available for resolution demands, is to synchronize a digitizing analog phase locked loop to the clock itself. The clock resolution attained doubles for every power of two the synchronized VCO frequency rises beyond the clock frequency itself. Limitations of this approach beyond those implied by the clock's accuracy will arise primarily from loop error (potentially noise induced), though this error should be small because the clock's phase is very steady.

A second alternative, considered again in Appendix D, involves generating an analog ramp waveform which is synchronized to the clock, and closely amplitude controlled, so that A to D converting the value of the ramp provides an interpolation of time between zero crossings of the highest digital clock bit available. Clearly, the limitation of this type of approach beyond the clock's accuracy is the accuracy of the ramp waveform generator, and A to D converter circuitry. Ash documents a similar

approach to this one for a period measurement circuit, although his approach involved a phase ramp which was triggered by carrier zero crossings and strobed by the clock's edge. After noise considerations, this approach resulted in effective time measurement with greater than a gigaHertz resolution (less than 1 nanosecond error). Similar resolution should be feasible with a ramp waveform which is clock synchronized, with the advantage that a single ramp waveform could be strobed (A to D converted) by several instruments in a system. If the digital clock strobing readout is intended to be a shared resource, a clock synchronized ramp seems a more logical choice in this respect.

A final approach involves analog to digital converting the phase modulated sinusoid itself, but in such a way to improve clock strobing measurements, rather than do arc-sin computations. This approach assumes that the phase modulated signal is a sinusoid, rather than just a one bit phase modulated digital signal, but a sinusoid could be generated by filtering in either case. Suppose the waveform is sampled at edges of the synchronized digital signal (quantized zero crossing time). This sampling is depicted below.

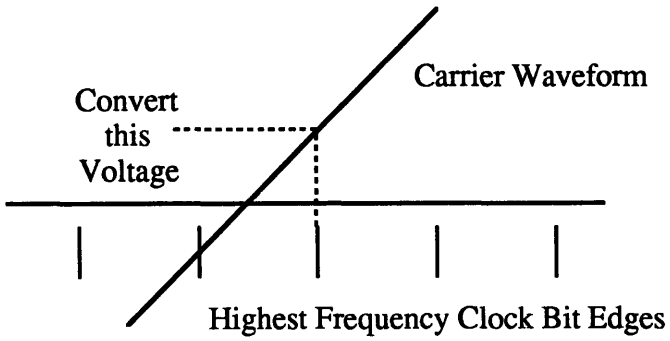

If there is no synchronization induced quantization error, then the value should be zero, but when a small quantization error exists, the A to D measurement is proportional to the quantization error and the carrier amplitude. The carrier amplitude can be measured by A to D converting the sinusoid at it's peak value, which occurs at a time

which can be readily estimated from the unadjusted quantized carrier zero-crossing times. The sinusoid peak value could be used to convert the amplitude-near-zero-crossing measurement into a phase measurement effectively, improving the resolution of the clock. This approach, suggested by Prof. Troxel of MIT, is analogous to the voltage ramp method previously discussed, with the voltage ramp being providing by the signal itself! This is a very good approximation near zero crossings, and the sinusoidal amplitude (and corresponding slope near zero crossings) can be determined very accurately each carrier half cycle without sampling at exactly the center of the peak, in the absence of noise

The errors which may result from this technique should originate from three primary sources. One error occurs because the peak measurement doesn't occur exactly at the sinusoid's peak, and the corresponding ramp slope will be underestimated. A second source of error results because the slope of the carrier sine wave is not truly constant around zero crossings. Both of these errors become smaller, as a percentage of total error, as the ratio between clock and carrier frequencies increases. If a nominal resolution of 10 bits, or a part in a thousand exists, then the worst case approximation error can be calculated from the worst case quantization errors yielding 15 bits effective improvement in clock resolution and 25 bits phase resolution if only these two errors are considered. The third major source of error is that induced by sampling and converting the voltages plus noise, as considered for the previous hybrid method.

Appendix C: Hybrid Phase Locked Loop Approaches

DPLL With Embedded Delta Sigma Modulation

Delta Sigma Modulators are frequently used in A to D and D to A conversion schemes because of their noise shaping properties. For this reason, it has been suggested that a digital phase locked loop employing a Delta Sigma modulator in the feedback loop might be a useful means of shaping quantization noise away from signal frequencies. A second reason why a delta sigma modulator may prove of use in a digital phase locked loop is that the resulting 1 bit digital signal could make NCO operation very simple, as it would only need to switch back and forth between two frequencies The Delta Sigma Phase Locked loop is pictured again below.

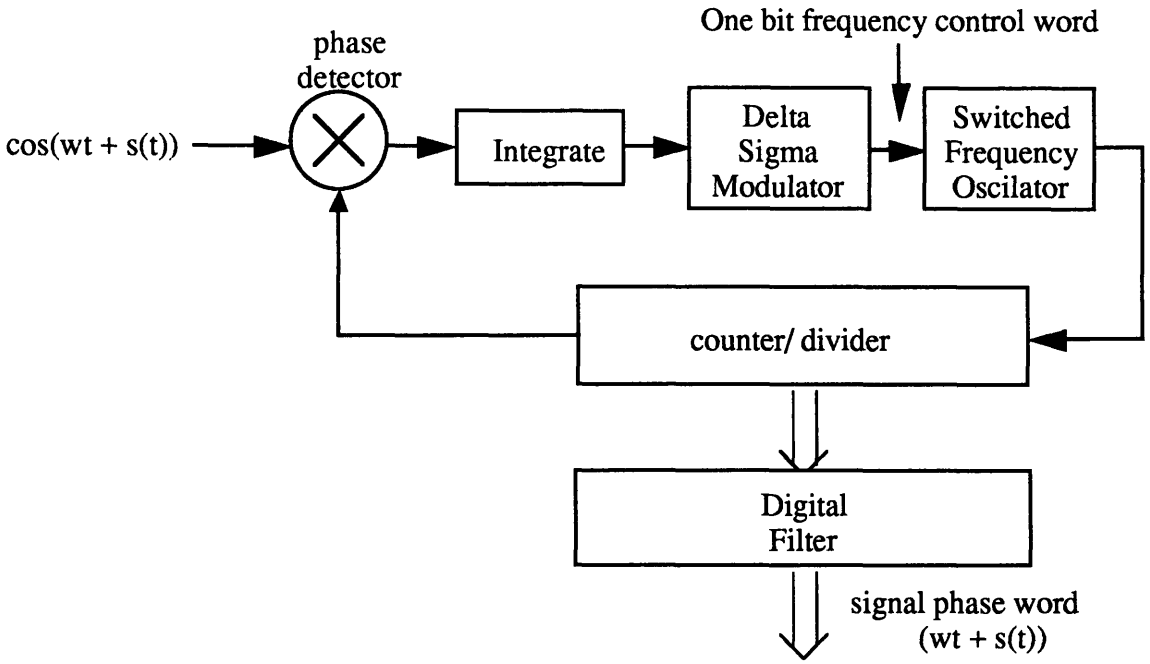

The Delta Sigma Phase Locked Loop

Delta Sigma Digital Phase Locked Loop with Digital Phase Detector

Based on the analysis and simulation of the Digital Clock Strobing Readout, and the following discussion of the fundamental resolution limitations resulting from synchronization induced quantization noise, it follows that no amount of noise shaping can improve resolution of the Delta Sigma Phase Locked loop beyond the limits described in Chapters 4 and 5. The potential benefits of the Delta Sigma Phase Locked loop can only be realized with an analog phase detector of some sort, which can produce a true null.

Delta Sigma DPLL with Analog Phase Detector

If the Delta Sigma phase locked loop contains an analog phase detector, the resulting phase data stream fundamentally changes. Consider the effects of a constant phase input

at an arbitrary nominal frequency. Still, each bit in the output of the Delta Sigma modulator, and each corresponding counter tick, corresponds to a high frequency clock bit of time. Phase errors corresponding to less than a clock bit of time take several carrier half cycles to resolve, but may eventually integrate (for sufficiently slow signals) to amount to a bit of phase. So rather than exhibiting a constant phase output, reported phase measurements vary with time but average to the correct value. In particular, since very small phase errors are detected and integrated, the resulting digitized phase output stream occasionally steps forward or backward in phase for a particular carrier half cycle to compensate for the integration of a phase error smaller than the clock's resolution which has existed for multiple cycles.

So in an average sense, the delta sigma phase locked loop keeps a perfect lock just as an ideal analog phase locked loop does. Unlike the analog loop, however, the resulting phase countdown in this case is synchronous with the clock, and no synchronization induced quantization error results because the counter's output is synchronously strobed. Hence for a constant (DC) phase input, the resolution limit of the Delta Sigma Phase Locked Loop is infinite, if an infinite number of phase samples are averaged.

For time varying signals, however, NCO induced quantization resolution measurements prevent infinite resolution by averaging, because only a finite number of phase measurements may be averaged. The dynamic situation is somewhat analogous to the realization of high resolution (for long beat cycles) when removing synchronization induced quantization noise discussed in the body of the thesis; patterns which emerge over many cycles can't be seen if the signal changes substantially before many cycles have been observed. In the case of Delta Sigma modulators, however, the effective resolutions which may be achieved have been well documented for A to D converter applications. Since the counter phase corresponds to delta sigma modulated version of the signal phase, with quantization of a clock bit of carrier phase, the same quantization

noise characteristics observed for a delta sigma analog to digital converter of the same order should be observed for a digital phase locked loop with a delta-sigma modulated phase output. With a second order delta sigma modulator, for example, signal to noise ratio improvements of approximately 10 dB per octave of oversampling ratio may be achieved[1]

For the Delta Sigma Phase Locked loop, the oversampling ratio is the ratio of carrier zero crossings to twice the phase modulating signal bandwidth. When the oversampling ratio is one, then the achievable phase measurement resolution for the Delta Sigma phase locked loop is no better than it would be for the any digital phase locked loop. For a large oversampling ratio, however, the resolution limitation should improve dramatically, regardless of the carrier and clock frequency values.

Hence using SFIR numbers once again, a 6400/50=128 oversampling ratio corresponds to 7 octaves which could should result in 70 dB or almost 12 bits of improvement beyond the resolution of a single phase measurement (with appropriate filtering of the NCO phase words). This is 12 bits of resolution beyond that which corresponds to the resolution of the clock, for a total of 23 bits of resolution.

Another Hybrid Phase Locked Loop

Another possible way to take advantage of the improved resolution feasible with an analog phase detector is to estimate phase more accurately than the time quantized demodulating signal, but correct for measured *phase differences*. The resulting phase difference can be predicted and subtracted to generate a *phase error* measurement, corresponding to the error in the NCO's estimation. This approach was partially explored by Sumner Brown, in conjunction with a Delta Sigma based loop A to D converter.

---

[1]Schreier reports improvements of 5 dB per octave per filter order up to a 40 dB per octave improvement in signal to noise ratio with oversampling ratio for an 8th order delta sigma modulator.

The idea of an NCO which is more accurate than the demodulating bit it generates seems difficult at first thought, but may be realized as depicted below.

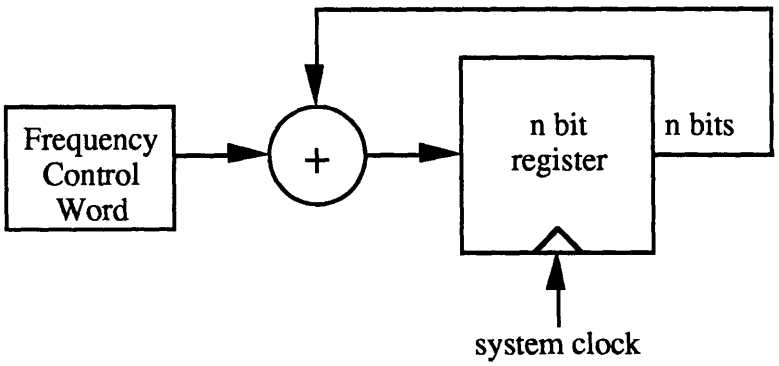

High Resolution Numerically Controlled Oscillator

The n bit register can be arbitrarily large, corresponding to increasingly large frequency control word values. If the frequency control word contains many bits, and changes in response to loop filtered error, then the n bit output may be used to represent n bits of carrier phase. The NCO phase jumps by an amount equal to the frequency control word value on each successive clock cycle. If the entire phase word of the NCO is to be phase locked with the carrier, and the MSB of the NCO is being fed back to generate a phase difference measurement, then the remaining bits of the NCO at the time when the MSB changes must be subtracted from the measured phase difference to determine the error between the NCO phase estimate and the carrier phase. Hence when the loop is in perfect lock, the phase difference measured by the phase detector corresponds to the NCO's n-1 LSB's at the time when the MSB flipped.

The resolution limiting factor of a hybrid loop using this configuration should be determined by the measurement accuracy of the phase detector. Since a physical null is not achieved at the phase detector, but instead predicted and subtracted, a mismeasured phase difference will result in a real phase error between the NCO and the carrier, hence a good linear phase detector is required to really benefit from this type of approach..

In either of the hybrid phase locked loops considered here, loop tracking error, as described for the analog phase locked loop, is still a potential source of error for a dynamically changing phase signal.

Appendix D: Universal Instrument Readout System

One potential application of interest for the clock strobing digital readout architecture is the development of a Universal Instrument Readout System (UIRS). The UIRS, in theory, is a central sensor signal digitization and demodulation front end which could be utilized in a variety of test, guidance, or other system applications where data from a variety of instruments needs to be converted to a multiplexed stream of data directed at a central computer.

UIRS Motivation

The motivations for such a Universal Instrument Readout System are numerous. A general instrument readout approach could allow greater flexibility in a guidance system by providing a common interface between an instrument and the guidance computer, increasing modularity within the system. Such a system would be easily upgraded when new technologies emerge, whether the new technology is a better instrument, faster signal processing, or changes in the guidance computer. Digital signal processing would allow new implementations with only software modifications. Also, a microprocessor based digital approach could multitask between different instruments within the system, eliminating redundancy and potentially lowering cost. In the lab, a

programmable general instrument readout could be connected to a data acquisition system and used for instrument testing.

Approach

A digital approach is demanded for the Universal Instrument Readout because ultimately, the information extracted from an instrument must be digitized in order to be fed to a computer. Also, the use of digital components provides greater flexibility to reconfigure the system for different instruments, and allow multitasking so that a single circuit can demodulate several instruments simultaneously. A high level block diagram the UIRS architecture is shown below.

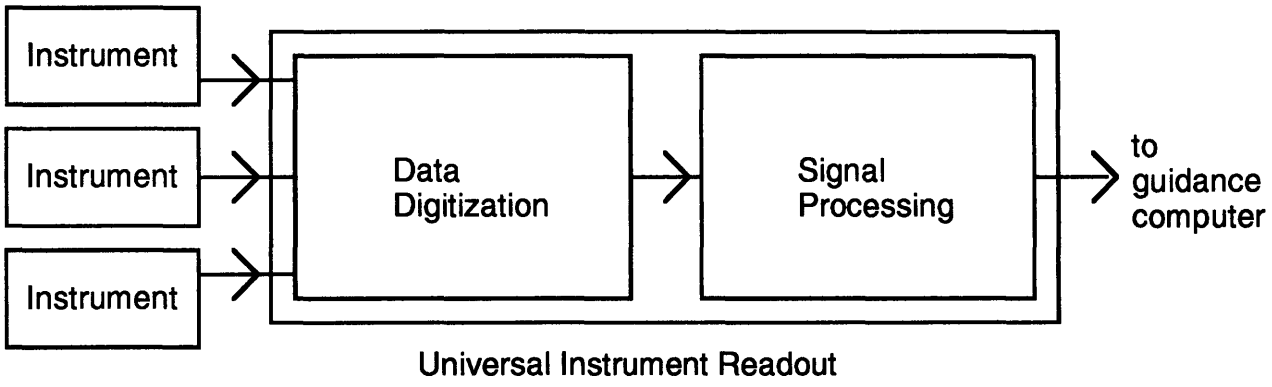

Data digitization in the UIRS would consist of two parts: a high resolution clock, for digitizing the *time of an event* and a highly accurate, low rate A to D converter, for digitizing *signal amplitudes*. Both of these subsystems would be universally available for instruments within the system. Each signal which requires event detection (i.e. any phase/freq. modulated signal) would require a dedicated event detector (comparator) and

holding register (for the strobed clock value). Signals which require A to D conversion would require a dedicated sample and hold circuit. These components, interfaced together, are shown below.

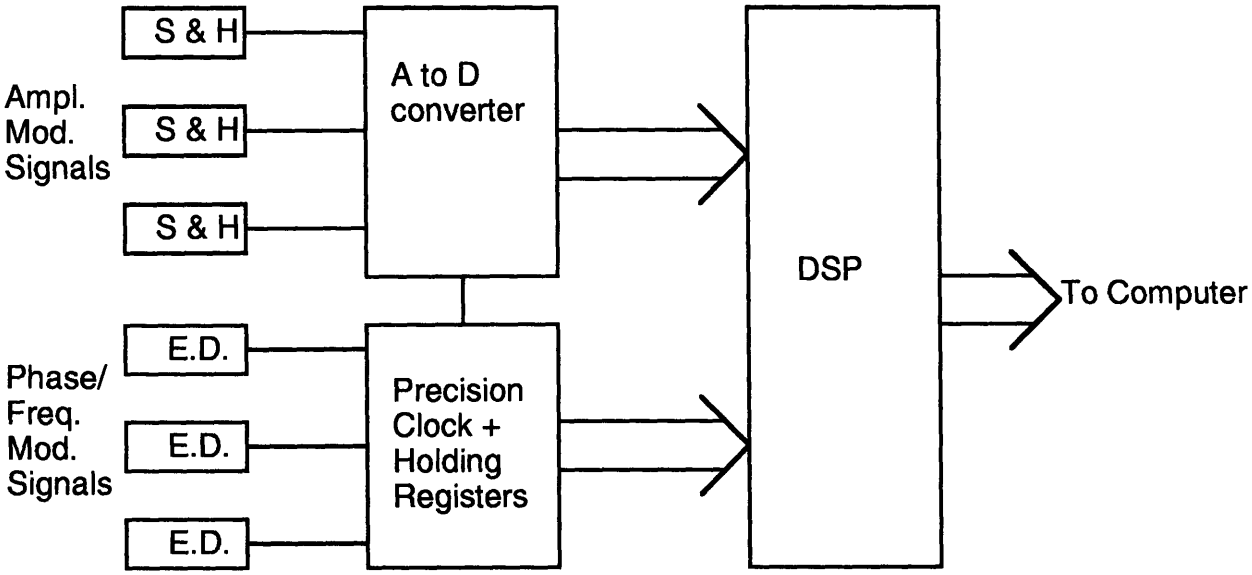

E.D. = Event Detector

A/D Conversion

Some instruments provide amplitude modulated signals, in which case zero crossing times do not hold enough information to retrieve the inertial information. In this case, an A to D converter is needed to provide the necessary amplitude information for conversion. Once again, the information is still a band limited signal, so sampling rates may be low for measurements of AM or unmodulated signals.

Very accurate, low rate A to D conversion could potentially be provided in two steps. First, an input voltage could be converted to a time modulated signal via a dual slope A to D, a voltage to frequency converter, pulse width modulator or similar voltage to time converter (V to T). These signals could be used in conjunction with the high

resolution clock to yield digital words for the DSP to process. Since all measurements taken would go the computer at fixed rates, dead time (when no values are being demanded at the DSP's output) could be used for A to D or V to F converter calibration, maintaining steady resolution of the analog measurements.

As with the high resolution clock and the DSP, the A to D converter could be a shared commodity. Each signal would only demand it's own sample and hold circuit, and the A to D converter could be able to access any of the sampled voltages at the proper time for digitization.

Achieving Fine Clock Resolution

Some instruments would only require a clock of several megahertz to provide adequate phase information from strobed clock readings. Other instruments would require higher rates. Equivalent rates of over a GHz could be achieved with several different approaches. One approach is to synchronize a high speed VCO to a lower stable frequency to provide extra bits of clock resolution. A second approach might be to synchronize a stable voltage ramp to a MHz clock bit and A to D convert the voltage to generate an interpolated time word. If a precision A to D converter is already available within the readout , achieving very high clock resolution would only require generating the ramp waveform and including sample and hold circuits at the event detector outputs.

Interestingly enough, it may be possible to simultaneously use the A to D converter to improve clock resolution while using the clock to achieve high resolution A to D conversion. Clock resolution could be enhanced by V to T converting the synchronous ramp voltage. The corresponding frequency/time signal from the V to T converter could then be measured with the non interpolated time bits. For very accurate A to D conversion, the V to T output could be measured with the clock and interpolator,

and a second V to T performed on the ramp waveform. This time only the non interpolated bits would be used. Hence two A to D conversions would be performed to yield a very high resolution voltage measurement. This process could be iterated to obtain arbitrarily fine resolution for either amplitude or time/phase measurements.

Application

The general instrument readout described might be capable of reading any instrument with narrow band information. Functionality of the instrument readout could be programmed in software, and the system could multitask it's time to readout several instruments simultaneously. Adding a new instrument to the readout would only require some new code and a small bit of additional hardware (comparator, holding register, and/or sample and hold circuit). The readout described might be well suited for either guidance system or instrument testing applications.

Bibliography

Advanced Inertial Technologies, Volumes I-III, The Charles Stark Draper Laboratory, Technical Report AFAL-TR-73-124,, March 1974.

Michael E. Ash, Anthony S. Kourepenis, et al, "Vibrating Quartz Accelerometer Readout and Signal Processing Techniques for Test and Guidance System Applications," Draper Lab, CSDL-R-2373, Nov. 1991.

Mel Bazes and Roni Ashuri, "A Novel CMOS Digital Clock and Data Decoder," IEEE Journal of Solid-State Circuits, VOL. 27, NO 12, Dec. 1992, pp. 1934-1940.

Kenneth C. Bell, Duncan B. Cox, Jr., Peter V. Hanks, "Resolver Angle Encoding Using the Digital Phase-Locked Loop," Technical Paper P-270, The Charles Stark Draper Laboratory, March 1976.

Sumner Brown, "High Accuracy SFIR Resolver Processing with a Numerically Controlled Oscillator," Draper Lab, ETD:92:456, June 23, 1992.

H. Bussey, "Resolver Readouts Using Phase Lock Loops," The Charles Stark Draper Laboratory, Memo No 20E:206:78, October, 1978.

Robert H. Comstock, "A Technique for Recovering the State of an Interrupted Digital Clock, Massachusetts Institute of Technology, Masters Thesis, T-519, May, 1969.

John Cook, "Digital Clock Phase Shifter without a Phase Locked Loop," IEEE Trans. on Circuits and Systems-1: Fundamental Theory and Applications, VOL. 40, No 4, April 1993, pp. 278-283.

Duncan B. Cox, Koen van Opijnen, J. Scott Rhodes, Louis Sutro, "Hybrid Phase (Hypha) Processing and Data Conversion," The Charles Stark Draper Laboratory, Report E-2697, January, 1973.

Robert H. Fall, "A Digital Preprocessor for a Specific Force Integrating Receiver, Massachusetts Institute of Technology, Masters Thesis, T-584, June, 1973.

Kenneth Fertig, Duncan B. Cox, "A New High-Speed General Purpose Input/Output Mechanism with Real-Time Computing Capability," The Charles Stark Draper Laboratory, Report E-2175, August 1967.

Edmund R. Foster, Robert H. Comstock, "Digital IMUE Study," The Charles Stark Draper Laboratory, Report R-841, Sept. 1974.

R.H. Frazier, "High-Accuracy Multi-Speed Wire-Wound Electromagnetic Resolvers," The Charles Stark Draper Laboratory, Report E-2794, August 1973.

F.M. Gardner, Phase Lock Techniques, 2nd ed. New York: Wiley, 1979.

Gurnam S. Gill, Someshwar C Gupta, "First-Order Discrete Phase-Locked Loop With Applications to Demodulation of Angle-Modulated Carrier," IEEE Trans. Commun., vol. COM-25, pp. 454-462, June 1972.

Capt. Albert E. Knauf, "An Integration Tecnnique for Aircraft Navigation Sensors, Massachusetts Institute of Technology, Masters Thesis, T-506, August 1968.

Howard Lewis, "A Multi-Speed Angle Encoder Using a Multi-Speed Resolver," Instrumentation Laboratory, Massachusetts Institute of Technology, Report E-1217.

William C. Lindsey and Chak Ming Chie, "A Survey of Digital Phase-Locked Loops," Proc. IEEE, vol. 69, pp. 410-431, Apr. 1981.

F.E. Natali, "Accurate Digital Detection of Angle-Modulated Signals," IEEE Electronics and Aerospace Systems Convention Record, 1968, pp. 407-413.

Holly C. Osborne, "Stability Analysis of an N'th Power Digital Phase-Locked Loop," IEEE Trans. Commun., vol. COM-28, pp. 1343-1364

Stanley S. W. Park, "Analysis of the Performance of a nonlinear Time Multiplexed, Noisy Channel in Transmitting Phase Encode Information, Massachusetts Institute of Technology, Masters Thesis, T-522, Sept. 1969.

Richard Schreier, "An Empirical Study of High-Order Single-Bit Delta-Sigma Modulators," IEEE Transactions on Circuits And Systems-II: Analog and Digital Signal Processing, VOL. 40, NO. 8, Aug. 1993.

Robert C. Tausworthe, "Theory and Practical Design of Phase-Locked Receiver, Volume 1, Technical Report No. 32-819, National Aeronautics and Space Administration, Jet Propulsion Laboratory, California Institute of Technology, Feb. 15, 1966.

Aaron Weinberg, Bede Liu, "Discrete Time Analysis of Nonuniform Sampling First- and Second-Order Digital Phase Lock Loops,", IEEE Trans. Commun., vol. COM-22, pp. 123-137, Feb. 1974.

Peter Welling, "Spectral Analysis of Delta Modulated and Pulse Width Modulated Waveform," The Charles Stark Draper Laboratory, Memo 83-15G-079, Feb. 18, 1983.