

A VLSI Demodulator for Digital RF Network Applications: Theory and Results

GARY J. SAULNIER, MEMBER, IEEE, CHARLES MCD. PUCKETTE, IV, SENIOR MEMBER, IEEE,
RICHARD C. GAUS, JR., MEMBER, IEEE, ROBERT J. DUNKI-JACOBS, AND
TIMOTHY E. THIEL

Abstract—An all-digital demodulator/detector is presented which is suitable for both analog FM and digital phase/frequency modulations. The system uses complex sampling, which employs a single A/D to sample the signal at an intermediate frequency (IF) and produce baseband in-phase (I) and quadrature phase (Q) signals, and a simplified technique for reducing the effect of the I/Q timing misalignment usually associated with this approach. The system also includes two detectors which operate simultaneously to provide noncoherent and differentially coherent detection, as well as automatic gain control (AGC) and automatic frequency control (AFC). The flexibility afforded by the two concurrent detectors in this all-digital system is shown to make it suitable for a wide range of applications. This paper describes the theory behind the demodulator/detector system and presents an implementation using a $1.25\text{ }\mu\text{m}$ bulk CMOS VLSI process. Methods are shown for extending and improving the I/Q sampling misalignment correction technique, as well as reducing the A/D sampling rate for a given IF frequency. Simulation and experimental results are shown to illustrate system performance for both analog and digital modulations.

I. INTRODUCTION

SECURITY and increased system capacity requirements, as well as the evolution of data-intensive applications in land mobile radio systems, have motivated the development of digital mobile radio systems. The desire to implement the digital receivers required by these new systems in a small, cost-effective manner has led to the development of concepts and architectures which maximize the use of VLSI technology. High-performance receivers, such as will be required in the second-generation digital cellular radiotelephone system now being developed, require a double conversion architecture.

Whereas the double conversion architecture has heretofore limited the use of VLSI technology to premodulation and postdetection processing such as speech encoders/decoders, equalizers, channel coders, etc., one would like the analog/digital interface to be at the antenna with all receiver functions performed digitally. Initial efforts in this direction have been previously reported [1].

The analog/digital interface can be moved closer to the antenna by using digital complex sampling techniques to

simultaneously sample an intermediate-frequency (IF) signal and then frequency-translate that signal to baseband [2]–[4]. Here, the IF signal is typically sampled at approximately four times the IF carrier and the samples are then sorted into the in-phase and quadrature-phase baseband signal channels. Coherent signal processing techniques may then be used to implement the required processes.

This paper discusses a receiver implementation that has been developed for detecting both analog FM and digital modulation formats such as FSK, CPFSK, and DPSK [5]. The input signal is sampled at IF by a single A/D converter and all the subsequent processing is performed digitally. Two detectors, which operate concurrently, are provided. These two detectors can be configured to provide frequency-to-voltage conversion, differentially coherent phase detection, envelope detection for automatic gain control (AGC), and frequency detection for automatic frequency control (AFC). A simplified technique for mitigating the effects of quadrature sampling misalignment is described. The detection subsystem has been implemented using a $1.25\text{ }\mu\text{m}$ CMOS process developed at GE. The methodology used for the design is presented and simulated demodulator performance is compared to experimental data. Potential extensions to the current design are identified and discussed.

II. DIGITAL DEMODULATOR

A block diagram of the general receiver architecture is shown in Fig. 1. An IF signal is fed into the demodulation subsystem after it has been frequency-translated from RF and bandlimited to reject adjacent channel interference. Focusing on the demodulation process, several functions can be defined: complex sampling and frequency translation to baseband, detection, postdetection processing, and ancillary functions, such as AGC and AFC. This section discusses each of these receiver functions.

A. Digital Complex Sampling

Digital complex sampling is a technique which uses IF sampling to produce baseband I and Q signals. Fig. 2 shows the classical analog approach and the corresponding digital approach. In the analog approach, Fig. 2(a), the IF signal is mixed with quadrature sinusoids at ap-

Manuscript received August 4, 1989; revised January 24, 1990.

G. J. Saulnier is with the Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180-3590.

C. M. Puckette, IV, R. C. Gaus, Jr., R. J. Dunki-Jacobs, and T. E. Thiel are with the Corporate Research and Development Center, General Electric Company, Schenectady, NY 12301.

IEEE Log Number 9036213.

0733-8716/90/1000-1500\$01.00 © 1990 IEEE




Fig. 1. Digital receiver block diagram.




Fig. 2. Coherent IF to baseband processes: (a) analog and (b) digital.

proximately the IF center frequency ω_{IF} . The outputs of the mixers are then low-pass filtered to remove the double-frequency terms and the resultant signals are sampled and quantized to develop the desired I and Q sample streams.

Fig. 2(b) shows a simple complex sampler. The IF signal is sampled at approximately four times the IF center frequency $4\omega_{IF}$ and the samples are sorted into the I and Q data streams. Fig. 3 shows the quadrature sinusoids used in Fig. 2(a) to frequency-translate the IF signal to baseband. The digital complex sampler, in effect, samples the input IF signal at times which correspond to peaks and zero crossings of these sinusoids. At sampling instant 1 ($t = 0$), the input signal is multiplied by unity in the I channel and zero in the Q channel. Therefore, the sample is passed directly to the I channel. At sampling instant 2 ($t = T/4$), the input signal is multiplied by unity in the Q channel and zero in the I channel, meaning that this sample is passed directly to the Q channel. At sampling instant 3 ($t = T/2$), the signal is multiplied by minus one (inverted) and passed to the I channel while, at sampling instant 4 ($t = 3T/4$), the signal is inverted and passed to the Q channel. This sequence is constantly repeating to produce the continuous I and Q streams.

A frequency decimation by two is performed in each of the I and Q streams by the removal of the zero-valued samples. This process results in the removal of the dou-




Fig. 3. Quadrature mixing signals.

ble-frequency terms which were present in the analog implementation of Fig. 2(a), since they now coincide with the sampling images. Low-pass filters are, therefore, not required. The price of this simplification, however, is that the I and Q samples are obtained at different time instants. This means that there is a time misalignment between the I and Q samples which, ideally, should be taken at the same instant in time.

The effect of the timing misalignment can be minimized by having a large ratio between the IF frequency and the signal bandwidth, so that the I and Q signals change slowly between samples. A consequence of having a high IF frequency, however, is that the A/D converter and digital circuits must operate faster. Therefore, increasing the IF frequency to reduce the effects of the sequential sampling has limited applicability. Other approaches, based on digital filtering, have been developed for reducing the misalignment at the expense of increased complexity [4].

A simple way to mitigate the effect of sequential sampling is to use an interpolation circuit to estimate the desired I or Q sample. The sequence of samples after the necessary inversions is

$$I_1, Q_1, I_2, Q_2, I_3, Q_3, \dots$$

If I_1 is paired with Q_1 , the aforementioned misalignment will result. However, if I_1 and I_2 are averaged, an estimate of the I signal at the time instant that Q_1 is taken will be obtained. Similarly, Q_1 and Q_2 can be used to produce an estimate of the Q signal at the instant I_2 is obtained. This simple two-point interpolation has proven quite effective at reducing the effect of the time misalignment when coupled with a large IF frequency-to-signal bandwidth ratio, i.e., a high degree of oversampling.

The digital demodulator incorporates the interpolator discussed above, along with a second averaging-type cir-

cuit to remove its sensitivity to dc components in the input signal or introduced by the A/D. Recall that the digital complex sampling process involves the inversion of some samples before they are placed in either the I or Q sample streams. For this process to be performed properly, the input signal must not have any dc bias. Before the inversions are performed, the sample stream is

$$I_1, Q_1 - I_2 - Q_2, I_3, Q_3, -I_4, -Q_4, \dots$$

If the signal contains a dc bias, a constant will be added to each of these samples. Once the inversions are performed, the dc bias on the inverted samples will be the negative of the original dc bias. Consequently, the I sample produced by inverting an $-I$ sample will have the inverted dc bias. The bias can be cancelled, then, by simply averaging an I sample that is obtained without an inversion and one that is obtained using an inversion. The same process is used for the Q samples. Naturally, this process assumes that there is sufficient oversampling to prevent the averaging process from distorting the signal.

Fig. 4 is a block diagram of the complex sampler, which incorporates both the misalignment reduction interpolator and the dc bias removal averager. Although the misalignment interpolator can be implemented on both the I and Q streams, it is sufficient to implement it in only one of the streams, as is the case in Fig. 4.

The above discussion provides a heuristic argument for the use of interpolation to mitigate the effects of I/Q misalignment. The same structure can be derived by considering the generation of an analytic signal using the discrete Hilbert transform (DHT). This approach provides a mathematical basis for the structure used in the digital demodulator and, additionally, indicates how the system can be extended and improved. An analytic signal is a complex function of a real variable whose real and imaginary parts are Hilbert transform (HT) pairs. Since the analytic signal is limited to positive frequencies, its use facilitates the processing of bandpass signals. Furthermore, since a real signal may be represented as the real part of the complex envelope of an analytic signal, means to implement an HT are of interest for the demodulation of bandpass signals. In practice, the HT is usually realized with an approximation to the DHT.

The generation of an analytic signal using a DHT may be accomplished in the manner outlined in Fig. 5. The bandpass signal [Fig. 5(a)] is sampled at a rate which satisfies the sampling theorem, thereby yielding the spectrum shown in Fig. 5(b). Passing this signal through the DHT filter network [Fig. 5(c)] will (ideally) yield the sampled analytic signal spectrum shown in Fig. 5(e). The network may be designed by first selecting the sampling rate, such that the negative sidebands of the sampled signal fall at $-j$ on the z plane and the desired positive frequency sidebands fall at $+j$. The filter is then synthesized so as to yield an N th-order zero at $-j$. An N th-order pole must be located at $z = 0$ for realizability. Fig. 5(d) illustrates this pole-zero configuration. Assume, for example,




Fig. 4. Complex sampler block diagram.




Fig. 5. Generation of analytic signals: (a) bandpass signal spectrum; (b) sampled bandpass signal spectrum; (c) DHT network; (d) z -plane pole-zero diagram; and (e) sampled analytic signal.

that $N = 2$. The desired z -plane transfer function is

$$H(z) = \frac{(z + j)^2}{z^2} \quad (1)$$

or

$$H(z) = -z^{-2} + j2z^{-1} + 1. \quad (2)$$

The realization of this network is shown in Fig. 6(a). A third-order network, that is obtained through the same process, is shown in Fig. 6(b). The third-order network has a wider bandwidth than the second-order network, meaning that improved I/Q misalignment correction can be obtained with less oversampling.

The required IF sampling process may, therefore, be implemented with a DHT by using the DHT network to remove the negative frequency components of the sampled signal and translating the resulting analytic signal to zero frequency. This process is depicted in Fig. 7. Because the filter network used to approximate the DHT is a linear network, the filtering and frequency translation processes may be interchanged as long as the appropriate sign changes are made in the filter network (see Fig. 8).

Embedding the network shown in Fig. 6 into the digital demodulator architecture, shown in Fig. 8, yields the result shown in Fig. 9(a); a first-order zero has been added at dc to remove dc offsets in the manner previously described. This system is functionally the same as the interpolator system described earlier. Fig. 9(b) shows a third-




Fig. 6. Discrete Hilbert transform networks: (a) second-order and (b) third-order.




Fig. 7. DHT digital sampler functional block diagram.




Fig. 8. Transposed DHT digital sampler functional block diagram.

order digital demodulator which will have a wider signal bandwidth than does the second-order demodulator [Fig. 9(a)]. It should be recognized that the DHT network can be replaced by a higher order network, or by a stagger-tuned network, if more signal bandwidth is required. The DHT network may also be combined with a low-pass filter FIR structure in the manner described by Parks, McClellan, and Rabiner [9].

B. Subharmonic Sampling

The A/D converter sampling rate limits the IF frequency that may be used with the digital demodulator described above. Receiver cost will also increase if the selected IF frequency requires a state-of-the-art A/D converter. A method that permits the relationship between the sampling rate and the IF frequency to be varied from the fixed 4:1 ratio required by the quadrature sampling process would be advantageous if it permitted higher IF fre-




Fig. 9. Complex digital sampler comparison: (a) second-order DHT and (b) third-order DHT.

quencies to be used with slower A/D converters. This section focuses on the use of subharmonic sampling, which allows use of sampling rates that are less than $4\omega_{IF}$ with the current demodulator implementation.

The sampling theorem that is used to select the A/D sampling rate is commonly stated as

$$F_s \geq 2F_H \quad (3)$$

where F_s is the sampling rate and F_H is the highest frequency of interest in the signal to be sampled. For a band-pass signal, this may be generalized to

$$\frac{2F_L}{K} \geq F_s \geq \frac{2F_H}{(K+1)} \quad (4)$$

where F_L is the lowest frequency of interest in the signal to be sampled and K is a nonnegative integer in the range of 0 to N . N , the number of permissible sampling frequency bands, can be calculated from

$$N = \left\lfloor \frac{F_H}{W} \right\rfloor \quad (5)$$

where $\lfloor x \rfloor$ is the floor function (i.e., the largest integer less than or equal to x) and

$$W = F_H - F_L. \quad (6)$$

For example, if $F_H = 4.1$ MHz and $F_L = 3.9$ MHz, $N = 21$. In other words, there are 21 different ranges of sampling frequency that can be used without spectral aliasing.

The quadrature sampling process described earlier used a sampling frequency that is $4F_{IF}$, which corresponds to $K = 0$. Since $F_s = 4F_{IF}$, the time between samples is $T/4$ where $T = IF$ signal period. In other words, the signal is sampled every 90° of the IF carrier frequency. It can be shown that the desired digital complex sampling process may be achieved with the sampling period generalized from $1/4F_{IF}$ to $(2n-1)/4F_{IF}$, $n = 1, 2, 3, \dots, N$. As a result, the sampling rate may be reduced for a given IF frequency or, conversely, the IF frequency increased for a given sampling rate.

C. Digital Detection Architecture

The demodulator structure was designed for applications involving several modulation schemes and requiring different detectors with multiple data rates. Data-rate flexibility is inherent in the scaling afforded by changing the clock rate. The detection requirements are met by including two detectors, which can be used to provide both non-coherent and differentially coherent detection of phase/frequency modulated signals. Although the two detectors can be implemented by reconfiguring one set of system functions, it is advantageous to operate both detectors concurrently. This aspect of the demodulator is discussed in a later section.

Fig. 10 shows the block diagrams of the detectors used in the demodulator. One of the detectors multiplies an undelayed I sample, $I(kT)$, by a delayed Q sample, $Q[(k - n_1)T]$, and an undelayed Q sample, $Q(kT)$, by a delayed I sample, $I[(k - n_1)T]$, where n_1 is the delay in samples and T is the sampling period. The output of the detector is the difference between these two products, i.e.,

$$s_1(kT) = Q(kT)I[(k - n_1)T] - I(kT)Q[(k - n_1)T]. \quad (7)$$

The second detector multiplies an I sample by a delayed I sample and a Q sample by a delayed Q sample. The output is the sum of the two products

$$s_2(kT) = I(kT)I[(k - n_2)T] + Q(kT)Q[(k - n_2)T] \quad (8)$$

where n_2 is the delay in samples.

Consider an input signal that consists of a generalized phase-modulated waveform

$$r(kT) = A(kT) \cos [\omega_{IF}kT + \Phi(kT)] \quad (9)$$

where T is the sampling period, $A(kT)$ is the sampled envelope, ω_{IF} is the IF frequency, and $\Phi(kT)$ is the information-bearing phase. Assume, for now, that $A(kT)$ is constant and equal to unity. The baseband I and Q signals are

$$I(kT) = \cos [\omega_{\Delta}kT + \Phi(kT)] \quad \text{and} \quad (10)$$

$$Q(kT) = \sin [\omega_{\Delta}kT + \Phi(kT)], \quad (11)$$

respectively, where ω_{Δ} is the difference between the local oscillator frequency (typically one quarter the sampling rate) and the IF frequency. The output of the first detector for these input signals is

$$s_1(kT) = \sin [\omega_{\Delta}n_1T + \Phi(kT) - \Phi((k - n_1)T)] \quad (12)$$

while, for the same inputs, the output of the second detector is

$$s_2(kT) = \cos [\omega_{\Delta}n_2T + \Phi(kT) - \Phi((k - n_2)T)]. \quad (13)$$




Fig. 10. Detector block diagrams: (a) sine detector and (b) cosine detector.

For obvious reasons, the first detector is called a sine detector while the second is called a cosine detector.

The argument of both the sine and cosine function, represented by (12), (13), consists of a term determined by the product of the frequency offset and the time delay, as well as a term which equals the difference in the phase of the delayed and undelayed signals $\Phi(kT) - \Phi[(k - n_i)T]$, $i = 1, 2$. The term dependent on the frequency offset usually can be considered a constant, since it is determined by local oscillator offset. The second term, however, provides information about the phase changes in the input signal.

For the case where the input signal is a frequency-modulated carrier

$$\Phi(kT) = 2\pi f_d \int_{-\infty}^{kT} x(\gamma) d\gamma, \quad (14)$$

where f_d is the frequency deviation constant and $x(\gamma)$ is the modulating signal. For n_1T small, the derivative of $\Phi(kT)$ can be approximated by

$$d/dt [\Phi(kT)] \cong \{\Phi(kT) - \Phi[(k - n_1)T]\} / n_1T \quad (15)$$

which is $1/n_1T$ times the phase-dependent term in both the sine or cosine detector outputs. From (14)

$$d/dt [\Phi(kT)] = 2\pi f_d x(kT). \quad (16)$$

From (15) and (16)

$$\begin{aligned} x(kT) &= d/dt [\Phi(kT)] / 2\pi f_d \\ &\cong \{\Phi(kT) - \Phi[(k - n_1)T]\} / 2\pi f_d n_1T. \end{aligned} \quad (17)$$

Therefore,

$$\Phi(kT) - \Phi[(k - n_1)T] \cong (2\pi f_d n_1T) x(kT). \quad (18)$$

Since $\sin \Theta \cong \Theta$ for small Θ , the output of the sine detector will be an amplitude-scaled version of the modulating waveform when $\Phi(kT) - \Phi[(k - n_1)T]$ and $\omega_{\Delta}n_1T$ are small. The simulated performance of the demodulator for FM tone modulation is described in Section III.

The sine detector is also useful for digital frequency modulations like frequency shift keying (FSK) and continuous-phase frequency shift keying (CPFSK) [5]. Included in this class of modulations are minimum shift keying (MSK) and Gaussian-shaped minimum shift keying (GMSK). The use of the sine detector with a delay duration that is less than one data bit duration provides noncoherent detection of these modulations. The action of the detector is to perform a frequency-to-voltage conversion.

In addition, both the sine and cosine detectors perform the necessary phase comparison to function as a differentially coherent detector, since they each contain a phase difference term. Here, a delay equal to an integer number of data bits is used in the detector. The choice of the sine or cosine configuration is dependent on the phase differences that are expected. The sine detector arrangement should be used if phase differences of 90° will occur, since the cosine detector will have a zero output for these differences. Similarly, the cosine detector provides nonzero output for phase differences of 0 and 180° .

As an example, the sine configuration should be used for differential detection of GMSK (with $h = 0.5$) with a 1 bit duration delay, while the cosine detector should be used for the same modulation with a 2 bit duration delay. In the first case, phase changes of 90° are expected while, in the second case, phase changes of 0 and 180° are expected.

D. Linearizer

The sine detector discussed above has a sinusoidal-shaped frequency-to-voltage conversion characteristic. For analog FM detection, the effect of the nonlinear characteristic is to introduce harmonic distortion into the detector output. Therefore, if the detector is to be used for analog FM in an application requiring low distortion, it is necessary to linearize the frequency-to-voltage characteristic.

The linearization process involves mapping the output voltage using an inverse sine function. Fig. 11 illustrates the frequency-to-voltage conversion characteristic of the sine detector and the desired linearized characteristic. The process of linearization can be accomplished by using either a read-only memory (ROM) or a polynomial approximation of the inverse sine function. In either case, the amplitude of the input IF signal must be known for the linearizer to operate properly, since the inverse sine operation requires an accurate mapping between output voltages and the peaks of the frequency-to-voltage characteristic. These voltages are determined by the input IF signal amplitude and are independent of the delay used in the detector.

E. Automatic Gain Control and Automatic Frequency Control

Generally, the IF signal in an FM receiver will come from a bandpass limiter, which ideally provides a constant amplitude signal. However, the actual output am-




Fig. 11. Ideal discriminator curves: (a) before linearization and (b) after linearization.

plitude from the bandpass limiter can vary with temperature and RF signal level. Although these variations may be relatively small, they can seriously degrade linearizer performance. As a result, it is necessary to incorporate some type of AGC circuit into the receiver. To provide AGC, the amplitude of the input signal must be constantly measured and any changes corrected. For constant envelope-type modulations, such as FM, the amplitude $A(kT)$ is a constant, i.e.,

$$A(kT) = \sqrt{I(kT)^2 + Q(kT)^2} = \text{constant}. \quad (19)$$

When the cosine detector is used with a delay of zero, i.e., n_2 equals zero, the output is simply the sum of the squares of $I(kT)$ and $Q(kT)$. Since the sine and cosine detectors are implemented in parallel, the cosine detector can be used for sensing amplitude while the sine detector is being used to demodulate the analog FM or detect the digital information.

This amplitude information can be used in a number of different ways to provide AGC. If fast amplitude variations were expected, it would most likely be necessary to use the information in a feedforward manner. However, since only slow variations are expected when a limiter is used, a feedback configuration should suffice. In the current configuration, the demodulator outputs a signal indicating that the input amplitude is either too low or too high. This information can be fed back to an AGC amplifier placed in front of the A/D converter. Fig. 22 (Section IV) shows the demodulator chip configured for AGC operation. Although the AGC function could be per-

formed using only digital processing by making the amplitude correction after the A/D conversion, performing the correction before the A/D converter has the advantage of matching the analog signal amplitude to the A/D converter.

Differentially coherent detectors display sensitivity to carrier frequency offsets, since these offsets affect the phase differences that are being measured. For some modulation schemes, even small frequency offsets can produce a severe reduction in detection performance. In these cases, it is necessary to provide some type of AFC circuit to minimize carrier frequency offsets.

Recall that the sine detector output contains a term that is determined by the product of the difference between the IF and local oscillator frequencies and the delay duration (12). Consequently, the sine detector output will be centered on a dc signal level that is determined by the frequency offset. If the frequency offset is removed, this dc signal will disappear. However, the detector output will also contain information due to the phase changes in the modulated signal, making it necessary to process the detector output in order to extract the desired information.

If the data-dependent portion of the sine detector output has zero dc value, a simple low-pass filter will provide the desired AFC control signal. Binary modulation with equally likely 1's and 0's would meet this requirement. More generally, however, it is necessary to obtain the median value of the detector output to avoid data dependence. The median can be found using two peak detectors and a subtractor, as illustrated in Fig. 12. The peak detectors determine the maximum and minimum output levels with the desired median value being midway between these values. The output of the frequency offset detector is low-pass filtered and fed back to a voltage-controlled oscillator in the IF or RF sections. Details of the AFC circuit operation can be found in [10].

Since the sine and cosine detectors operate concurrently, the cosine detector can be used for differentially coherent detection while the sine detector is used as part of the AFC circuitry. Similarly, the cosine detector can function as part of the AGC circuit when the sine detector is being used for high-fidelity analog FM.

III. SIMULATION RESULTS

Computer simulation of the demodulator was performed prior to chip fabrication in order to allow the identification and correction of any design problems. The simulation generated the output spectrum for the sine detector when the input is a tone-modulated FM carrier. This type of input signal was used because it produces an output spectrum which clearly indicates the presence of detector problems. The output spectrum was produced using three methods: inserting signal parameters into the equations for the detector output, performing floating-point and integer simulations of the detectors, and, finally, using the computer model of the designed chip. This last simulation is possible because the chip was designed using the General Electric nibble-serial silicon compiler [11], which




Fig. 12. AFC configuration.




Fig. 13. Theoretical sine detector output spectrum.

permits the designer to develop, as well as simulate, a pipelined DSP architecture.

The FM signal has a modulating tone of the form

$$x(kT) = \cos(\omega_m kT) \quad (20)$$

where ω_m is the angular frequency of the tone. Fig. 13 shows a plot of the power spectrum of $s_1(kT)$ (12) with $x(kT)$ as defined above. This plot only shows the spectrum resulting from the expression for $s_1(kT)$ and not of the output of the demodulator architecture. Fig. 13 was used as a baseline for comparison to the spectral plots of the output of the demodulator. All simulations employing tones used the following values for the parameters:

$$f_d = 75 \text{ kHz}, \quad n_1 = 8, \quad T = 125 \text{ ns},$$

$$f_m = 37 \text{ kHz}, \quad \text{and} \quad f_\Delta = 2 \text{ kHz}.$$

Both even and odd harmonics are present in $s_1(kT)$, as illustrated in Fig. 13. The presence of the even harmonics is due to the 2 kHz offset between the signal intermediate frequency and the local oscillator (sampler). If no frequency offset was present, only the odd harmonics would exist. This can be seen by expanding $s_1(kT)$ into a Taylor series for both cases. Note also that the magnitude of both even and odd harmonics decreases with increasing frequency. Recall that the harmonics are due to the sinusoidal frequency-to-voltage characteristic of the sine detector and can be removed by a linearizer.

Next, the sine detector architecture was simulated using floating-point arithmetic; the resulting output spectrum is shown in Fig. 14. Note that only the second and third




Fig. 14. Sine detector output spectrum from floating-point simulation.




Fig. 16. Expanded view of the spectrum of Fig. 15.




Fig. 15. Sine detector output spectrum from chip simulation.




Fig. 17. Sine detector output spectrum using 10 bit quantization.



harmonics are clearly visible and that their magnitudes, with respect to the fundamental, agree with the relative magnitudes shown in Fig. 13. The remaining harmonics have been lost in the noise floor associated with the output of this architecture. In comparing Figs. 13 and 14, any harmonic 80 dB or so below the fundamental will be lost in the noise and, therefore, only the second and third harmonics are visible.

Fig. 15 shows the spectrum of the output of the sine detector when the simulation was done using the chip simulator contained within the compiler. The assumed parameters are the same as those referred to above. Note that the harmonics have the same relative magnitudes as in the two previous cases, but now an unexpected effect has appeared: there are two sidebands surrounding each harmonic. Fig. 16 is an expanded view of Fig. 15, in which it can be seen that the sidebands are separated from their respective harmonics by ± 8 kHz. These spurious sidebands are most likely due to the combination of the frequency offset between the nominal intermediate frequency and the local oscillator (sampling rate), and the noise introduced by the 10 bit quantization. There would be a pair of sidebands around the fundamental if it were not for the operation of the I/Q time alignment interpolator, whose purpose it is to remove these sidebands.

To test the hypothesis that A/D quantization contributed to these undesired spectral components, an integer simulation was done for 10 bit and 6 bit quantization. Fig. 17 is the power spectrum of the output for 10 bit quantization, while Fig. 18 is for 6 bit quantization. Note the prominence of the sidebands as they dominate the two remaining harmonics. Clearly, as the coarseness of quantization of the intermediate frequency signal increases, the relative magnitudes of the sidebands increases. Thus, one would expect that, as quantization becomes finer, the sidebands would decrease. This is the case for the spectrum of the output of the floating-point simulation, which has little associated quantization noise.

The effect of the increase in the level of the sidebands can be explained qualitatively from the work of Park [6] for an analog architecture similar to that of the digital demodulator chip. His analysis showed that, for ideal detection, the output noise spectrum will have a shape that peaks in a manner similar to the sideband envelope observed in Figs. 15, 17, and 18.

As a final test, the chip simulator was used to determine the sine detector performance for Gaussian minimum shift keying (GMSK) [7] modulation. The detector was configured for 1 bit differential detection in an additive white-Gaussian noise channel and the results were compared to





the published results of Simon and Wang [8] for a similar normalized received bandwidth $B_r T$. Fig. 19 shows a plot of probability of bit error versus E_b/N_o for the simulation. For this simulation, $B_r T = 1.25$. A plot of the curve for $B_r T = 1.1$ from [8] is shown for comparison. In both cases, the normalized transmit bandwidth is $B_r T = 0.5$. As can be seen, the results for the sine detector are in good agreement with the results of Simon and Wang.

IV. INTEGRATED CIRCUIT IMPLEMENTATION

The direct digital demodulator IC, version 1 (DDM1) is a programmable digital demodulator containing the sine detector, AGC control (envelope detector), and an externally controllable integrate-and-dump circuit for sample rate decimation prior to signal output. The external control is provided to allow synchronization with a bit timing recovery (BTR) circuit for decision timing. A block diagram of DDM1 is shown in Fig. 20.

The DDM1 integrated circuit accepts input words of up to 10 bits at a word rate of up to 16 MHz. Fig. 20 shows two demodulator sections fed by a common sorting and offset correction circuit. Demodulator parameters, such as the sine detector delay, AGC setpoint, integrate-and-dump



durations, and various scaling factors, are set using a serial programming chain.

A microphotograph of the chip is shown in Fig. 21. The pipeline architecture used by the nibble-serial compiler is evident. Fig. 22 shows the DDMI interfaced to currently required off-chip components.

A. Experimental Measurements

The performance of the demodulator architecture was tested experimentally using a breadboard implementation. This breadboard incorporated the complex sampler and sine detector circuits. All processing was performed using 8 bit quantization. The performance of the averaging circuits in the complex sampler and the subharmonic sampling process were of particular interest.

Fig. 23 is a set of discriminator (i.e., frequency-to-voltage conversion) curves for the sine detector. For these




Fig. 22. Typical DDM1 application including AGC circuit.




Fig. 23. Discriminator curves for digital demodulator breadboard.

curves, the sampling rate is 16 MHz and the IF frequency is 4 MHz. The parameter is the detector delay, in measured units of the 8 MHz clock (one-half of the 16 MHz clock due to the sorting into *I* and *Q* streams) periods. The sinusoidal nature of the characteristics are evident.

Keeping the sampling rate constant, the IF frequency was increased to produce the subharmonic sampling scenario. Fig. 24 is a discriminator characteristic for $F_S = 4/3 F_{IF}$, meaning that $F_{IF} = 12$ MHz. Notice that the slope of the curve is opposite that for the $F_S = 4 F_{IF}$ condition. This change is due to the fact that this sampling arrangement effectively reverses the polarity of either the *I* or *Q* channel, thereby inverting the detector output. Fig. 25 is the result for $F_{IF} = 36$ MHz, which has $F_S = 4/9 F_{IF}$. The discriminator curve now has the same polarity as for $F_S = 4 F_{IF}$. The reduction in amplitude is due to the fact that the frequency of the input signal is now outside the power bandwidth of the A/D converter.

The performance of the *I/Q* misalignment correction and dc removal circuits was tested using a carrier that was FM-modulated by a tone. For the tests, $f_d = 75$ kHz, $f_m = 50$ kHz, $F_{IF} = 4$ MHz, $f_d = 5$ kHz, and the detector delay is seven samples. Fig. 26 shows the spectrum of the sine detector output in the region of the 50 kHz tone when






Fig. 24. Discriminator curve for $F_S = 4/3 F_{IF}$.Fig. 25. Discriminator curve for $F_S = 4/9 F_{IF}$.

Fig. 26. Output spectrum without correction circuits.

the *I/Q* misalignment and dc removal circuits are disabled. Two pairs of sidebands, one at ± 5 kHz and another at ± 10 kHz, appear around the desired 50 kHz tone. Due to the large degree of oversampling, i.e., the sampling rate is much larger than the bandpass bandwidth of the signal, these sidebands are approximately 50 dB below the desired tone.

Fig. 27. Output spectrum with I/Q misalignment correction.Fig. 28. Output spectrum with I/Q misalignment correction and dc removal circuit.

Introducing the I/Q misalignment correction produces the spectrum of Fig. 27. Clearly, the outer pair of sidebands have been greatly attenuated. Similarly, adding the dc removal circuit along with the I/Q misalignment correction attenuates the second pair of sidebands, producing the spectrum of Fig. 28.

V. CONCLUSIONS

The all-digital demodulator/detector system that has been presented is suitable for analog FM, as well as digital phase/frequency modulations. The dual detector structure is a major factor in providing this flexibility by allowing both noncoherent and differentially coherent detection processes. In addition, since both detectors run concurrently, one detector can provide information for either AFC or AGC operations while the other provides signal detection.

The combination of the subharmonic sampling and interpolator-based I/Q misalignment correction circuit makes it possible to interface at a high IF frequency without using excessive sampling rates. The actual sampling rate is determined by the signal bandwidth, allowing suf-

ficient oversampling to allow the I/Q misalignment interpolator to adequately mitigate any misalignment effects. The amount of oversampling that is "sufficient" is clearly a function of the particular modulation in use and the amount of degradation that is acceptable. In any event, the use of oversampling, in combination with the interpolator misalignment correction, is an attractive alternative to other approaches which minimize sampling rates at the expense of complicating the correction processing.

REFERENCES

- [1] I. A. W. Vance, "An integrated V.H.F. radio receiver," *Radio and Electron. Eng.*, vol. 50, pp. 158-164, 1980.
- [2] D. W. Rice and K. H. Wu, "Quadrature sampling with high dynamic range," *IEEE Trans. Aerospace Electron. Syst.*, vol. 18, pp. 736-739, Nov. 1982.
- [3] J. L. Brown, Jr., "First order sampling of bandpass signals—A new approach," *IEEE Trans. Inform. Theory*, vol. IT-12, pp. 613-615, Sept. 1980.
- [4] C. M. Rader, "A simple method for sampling in-phase and quadrature components," *IEEE Trans. Aerospace Electron. Syst.*, vol. AES-20, pp. 821-824, Nov. 1984.
- [5] J. B. Anderson, T. Aulin, and C.-E. Sundberg, *Digital Phase Modulation*. New York: Plenum, 1986.
- [6] J. H. Park, Jr., "An FM detector for low S/N," *IEEE Trans. Commun. Technol.*, vol. COM-18, no. 2, pp. 110-118, Apr. 1970.
- [7] K. Murota and K. Hirade, "GMSK modulation for digital mobile radio telephony," *IEEE Trans. Commun.*, vol. COM-29, no. 7, pp. 1044-1050, July 1981.
- [8] M. K. Simon and C. C. Wang, "Differential detection of Gaussian MSK in a mobile radio environment," *IEEE Trans. Vehicular Technol.*, vol. VT-33, no. 4, pp. 307-320, Nov. 1984.
- [9] J. H. McClellan, T. W. Parks, and L. R. Rabiner, "A computer program for designing optimum FIR linear phase digital filters," *IEEE Trans. Audio Electroacoust.*, vol. AE-21, pp. 506-526, Dec. 1973.
- [10] G. J. Saulnier and W. Rafferty, "DSP-based non-coherent dual detector demodulation for land mobile radio channels," in *Proc. Conf. Record IEEE Int. Conf. Commun.*, Toronto, Canada, June 1986, pp. 1008-1012.
- [11] R. Hartley and P. Corbett, "A digit-serial compiler operator library," in *Proc. Design Automat. Conf.*, 1988.



Gary J. Saulnier (S'80-M'84) was born in Fall River, MA, on November 25, 1958. He received the B.S., M.E., and Ph.D. degrees in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1980, 1982, and 1985, respectively.

In 1984, he joined the Communication Technology and Systems Branch of the General Electric Corporate Research and Development Center, Schenectady, NY, where he studied the design and implementation of bandwidth-efficient digital modulation techniques for fading channels. Since 1986, he has been a member of the Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute. His research interests include narrowband modems for mobile and mobile-satellite channels, antijam spread spectrum communications, and electronic instrumentation.

Dr. Saulnier is a member of Eta Kappa Nu.



Charles McD. Puckette, IV (S'57-M'62-SM'77) was born in NJ. He received the B.S.E.E. degree from Lafayette College, Easton, PA, in 1959, the M.S.E.E. degree from the University of Pennsylvania, Philadelphia, in 1961, and graduated from GE's Advanced Course in Engineering in 1967.

From 1959 to 1961, he worked on radar transmitters at the Missile and Surface Radar Division of RCA. Following 18 months of service in the U.S. Army Ordnance Corp as a First Lieutenant, he joined the Heavy Military Electronics Depart-

ment of the General Electric Company, where he was responsible for the analysis and design of radar receivers and related signal processing equipment. In 1967, he transferred to the General Electric Corporate Research and Development Center, Schenectady, NY, where he has worked in the areas of communication systems and signal processing. He was appointed Manager, Communication Programs, in 1983, and currently supervises a group of 25 people conducting research and development on digital cellular mobile radiotelephone systems and technology, SIGINT modeling and simulation, source data compression, and the development of interactive CAD environments for both communication system modeling and simulation, and digital signal processing systems. Prior to his present position, he was Manager, Communication Technology. In this position, he supervised and made technical contributions leading to the development of power-line carrier data networks, monolithic communication functions, fiber-optic data networks, and mobile radio telephone systems such as GE's personal radio communication system (PRCS). In addition to numerous internal General Electric reports, he has written or coauthored 28 technical papers and articles pertaining to his fields of work, has been awarded 27 U.S. patents, and has two patents pending.

Mr. Puckette is a member of Eta Kappa Nu and is a Licensed Professional Engineer in the State of New York. Active in professional society affairs, he was an Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS for three years, served as Chairman of the Signal Processing and Communication Electronics Technical Committee of the IEEE Communications Society for three years, and has organized and chaired several technical sessions at IEEE conferences.



Richard C. Gaus, Jr. (M'87) was born in Norfolk, VA, on November 11, 1951. He received the B.S. degree in physics from Millersville State College, Millersville, PA, in 1977, and the M.Eng. degree in electrical engineering from the University of Virginia, Charlottesville, in 1980.

In 1977, he worked at the General Electric Company, Waynesboro, VA, as a Cooperative Engineer in connection with his master's degree program at the University of Virginia. At the University of Virginia, he worked as both a Teaching

and Research Assistant during the course of his studies in communications engineering. Upon graduation in 1980, he began working at M/A-COM

Research Laboratories, Rockville, MD, in digital communications and signal processing, with emphasis on both speech and image data compression. Since 1984, he has been with the General Electric Corporate Research and Development Center, Schenectady, NY. His current interests are in the implementation of speech signal processing algorithms for data compression and in the design of digital communication systems for applications in portable and mobile radio.



Robert J. Dunki-Jacobs received the B.S. degree and the M.Eng. degree in nuclear engineering from Cornell University, Ithaca, NY, and the M.S. degree in electrical engineering from Union College, Albany, NY.

He is currently a Staff Member at the General Electric Corporate Research and Development Center, Schenectady, NY, where he has spent 15 years involved with computers, computer architectures, and communications. He is the designer of four custom integrated circuits for digital signal

processing applications. He is the author of four papers and holds three U.S. patents.



Timothy E. Thiel was born in Providence, RI, on February 4, 1964. He received the B.S. degree in computer and systems engineering, and the M.Eng. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1986 and 1988, respectively.

He has been at the General Electric Corporate Research and Development Center, Schenectady, NY, since 1987. While at GE, he has worked in the areas of FM communications, phased-array medical imaging, and the simulation of communications systems.