Digital Carrier Demodulation for the DSN Advanced Receiver

R. Sadr and W. J. Hurd
Communications Systems Research Section

The digital in-phase and quadrature (I&Q) carrier demodulation for the Deep Space Network's Advanced Receiver is described and analyzed. The system provides coherent demodulation for a variety of modulation formats including Binary Phase Shift Keying (BPSK), BPSK with carrier residual, Quadrature Phase Shift Keying (QPSK), Offset-QPSK (OQPSK), and Minimum Shift Keying (MSK). The article focuses on the theory and hardware design of the halfband filters which are the integral part of the demodulator. The underlying theory of the filters is summarized, a breadboard hardware design is described, and a VLSI implementation is proposed which significantly decreases the hardware. A second important problem analyzed is DC-offset in the demodulator. This is a serious problem which translates into bias error in the residual carrier phase detector. The dynamic range of the complex mixer is analyzed using a probabilistic approach. It is deduced that the resulting static phase error is less than 0.2 deg when the ratio of carrier power to noise power in the demodulator input bandwidth is −72 dB or higher. Thus, the static phase error is negligible at a carrier power to noise spectral density of 0 dB-Hz for a 15-MHz bandwidth demodulator.

I. Introduction

A new Advanced Receiver is presently under development for future implementation throughout the NASA Deep Space Network (DSN). The new receiving system will perform carrier and subcarrier demodulation, data detection, and Doppler extraction and will interface closely with the exciter, the ranging assembly, and the monitor and control subsystem. This article is concerned only with the digital carrier demodulator. We present the architecture and analysis of a digital demodulator for coherent carrier demodulation of the general class of amplitude and phase modulation. Besides the current standard DSN modulation type of binary phase shift keying (BPSK) on subcarriers, with residual carrier [1], the demodulator is appropriate for Quadrature Phase Shift Keying (QPSK), Minimum Shift Keying (MSK), Staggered or Offset QPSK (OQPSK), and the standard DSN modulation but with fully suppressed carrier.

The key characteristics of the demodulator are the passband width and ripple; the stopband width, attenuation, and ripple; and the DC-offset or bias introduced by the numerical accuracy effects. The DC-offset is critical because it causes bias in measurement of the phase of residual carriers. In this article, we specify the architecture of the demodulator, design

and analyze a practical implementation of the arm filters as halfband filters, and determine the numerical accuracy required to meet the DSN's needs.

A. Previous Work

There are many references for a complete treatment of digital modulation techniques such as BPSK and QPSK [2], [3]. The equivalency of staggered QPSK (OQPSK) and MSK is shown in [3]; so-called serial MSK [3] is only equivalent to OQPSK when the input symbols at the transmitter are differentially encoded. BPSK with carrier residual is described in [1]; the architecture of the old version of the Advanced Receiver and tracking loops is discussed in [4]. The effects of offset in sampling and the performance of the integrate-and-dump filter as a matched filter are discussed in [5]. To reduce the loss incurred due to offset sampling, the weighted integrate-and-dump is proposed and analyzed in [6]. References [7] and [8] provide extensive treatments of digital filter design. The loss due to amplitude ripple within the passband of the anti-aliasing filter is discussed and analyzed in [9]. Originally, Bellanger [10], [11] introduced the halfband filters for computation and storage reduction in multi-rate digital filtering applications, and he further analyzed the filter order, multiplication rate, and storage requirement of the halfband filters. Reference [12] provides the hardware description and test results for the older version of the Advanced Receiver.

B. Outline of Article

The receiver structure is described in Section II. In Section III, the underlying theory of halfband filters is discussed. In IIIB, a design example is given and the hardware architecture of a single board that has been fabricated and demonstrated in the laboratory is described. A VLSI chip is outlined in IIIC that results in significant reduction in the chip count. In Section IV, DC-offsets resulting in phase bias are analyzed, including the effects of numerical precision in the table for sine and cosine generation, and rounding. A probabilistic approach is used. Finally, in Section V, conclusions and some final observations are presented.

II. Receiver Structure

A simplified block diagram of the receiver, showing portions of the receiver impacting demodulator design, is shown in Fig. 1. The received signal from the S-, X- or K-band low noise amplifier is open-loop converted from radio frequency (RF) to intermediate frequency (IF) in highly stable, wideband RF-to-IF downconverters, which now exist for S- and X-band. The downconverter output goes into the IF strip. In the IF strip, the carrier phase locked loop is closed by mixing the signal against the numerically controlled oscillator (NCO) output, resulting in a fixed second IF frequency. This signal is filtered by an anti-aliasing bandpass filter, and down-converted to a final IF frequency. The final IF signal is sampled by an A/D converter whose output goes to the digital demodulator. The demodulator outputs digital in-phase and quadrature components to the baseband signal processor, which measures all statistics needed by the software algorithms and outputs detected telemetry symbols to the decoder(s).

The final IF frequency has been chosen to be 10 MHz in order to be compatible with existing DSN systems, such as the ranging system. Thus, at the final IF output, the carrier is phase locked to 10 MHz. The passband is approximately 10 ± 7.5 MHz, as limited by the anti-aliasing bandpass filters at the previous IF frequency.

The demodulator block diagram is shown in Fig. 2. The input is 8-bit digitized IF, with the carrier phase locked to 10 MHz. The sampling rate is a fixed frequency near, but not exactly, 40 MHz. A tentative value is 39.6 MHz. The offset from 40 MHz is to minimize self-induced radio frequency interference (RFI) at the carrier frequency of 10 MHz. The demodulator multiplies the input data by sine and cosine waveforms at 10 MHz, resulting in in-phase and quadrature baseband signals plus sum frequency terms. The sum frequency terms are filtered out by lowpass "arm" filters.

Throughout this article we denote the sampling period as T_s, the sampling rate as $f_s = 1/T_s$, the symbol period as T, the IF frequency as f_c, and the radian IF frequency as $\omega_c (= 2\pi f_c)$.

The input samples are demodulated by the reference in-phase $\cos(\omega_c n T_s + \hat{\phi}_c)$ and quadrature components $\sin(\omega_c n T_s + \hat{\phi}_c)$, where $\hat{\phi}_c$ denotes the carrier phase reference. These are generated from a look-up table. Ideally, the multiplier output signals are filtered by a zonal lowpass filter, to filter out the double frequency images that result from the multiplication in the time domain of the signal by a sinusoid.

The original spectrum, the mixer output spectrum, and the desired filtered output spectrum are shown in Fig. 3. Since the lowpass bandwidth after demodulation and filtering is only half the IF bandwidth, or 7.5 MHz, the sampling rate is reduced by a factor of two, resulting in the final spectrum in Fig. 3. In this article, the application of halfband filters is considered for low-pass filtering of the demodulated signal. This class of filters is a special case of lowpass Finite Impulse Response (FIR) filters which is particularly appropriate for removing mixer sum frequency terms and for use with a decimation factor of two.

III. Halfband Filters

To meet the specification for a desired frequency response of a digital filter, it is desirable to minimize the order of the filter (denoted as N), thereby reducing the number of multiplication and addition operations required to implement the filter. We have selected the halfband filters to eliminate the double frequency images that are produced by complex demodulation. The following points are the main considerations that led us to select the halfband filters:

(1) *Reduced hardware.* Halfband filters require approximately half the number of multiplications and additions for a given N, when compared to an arbitrary linear phase digital filter. This is because almost half of the N filter coefficients are zero.

(2) *Equal passband and stopband width.* Stopband width is the same as passband width. This is ideal for eliminating sum frequency terms resulting in complex heterodyning, because the sum frequency terms to be eliminated have the same bandwidth as the difference frequency terms to be passed.

(3) *Equal passband and stopband ripple.* The ripple in the passband and the stopband is the same in peak deviation. This results in approximately the same magnitude of signal-to-noise ratio loss due to passband ripple [9] and due to sum frequency noise aliased into the passband after decimation by two.

(4) *Ease of implementation.* When a decimation factor of two is employed, the processing rate is also reduced by a factor of two.

For equiripple design of Finite Impulse Response (FIR) filters, it is necessary to specify a set of tolerance parameters in order to practically implement these filters. The FIR design problem is then formulated as a Chebyshev approximation method [7], [8]. These parameters, as shown in Fig. 4, are:

δ_p = ripple (deviation) in the passband from ideal response

δ_s = ripple (deviation) in the stopband from ideal response

$f_p = \omega_p/2$ = normalized passband edge frequency

$f_s = \omega_s/2$ = normalized stopband edge frequency

For the case when $\delta_p = \delta_s = \delta$ and $f_s = 0.5 - f_p$, then the resulting equiripple optimal solution to this approximation problem has the property that

$$H(e^{j\omega}) = 1 - H(e^{j(\pi-\omega)}) \tag{1}$$

that is, the frequency response of the optimal filter is odd-symmetric around $\omega = \pi/2$, and at $\omega = \pi/2$

$$H(e^{j\pi/2}) = 0.5 \tag{2}$$

It is shown in [8] that any symmetric FIR filter satisfying (1) has an impulse response of the form

$$h_k = \begin{cases} 1; & k = 0 \\ 0; & k = \pm 2, \pm 4, \pm 6, \ldots \end{cases} \tag{3}$$

Hence, all of the even coefficients are zero, except for $k = 0$. This reduces the complexity required to implement this class of filters. Particularly when a decimation factor of 2 is used, the required number of multiplications is half of that needed for symmetrical FIR designs and one-fourth of that for arbitrary FIR designs. The SNR loss due to amplitude ripple is studied in [9].

A. Halfband Architecture with Decimation Factor of 2

In this section, we outline a simplified architecture for an Nth (N-odd) order halfband FIR filter as shown in Fig. 5(a). We denote the filter output as y_n, the input as x_n, and the filter coefficients as h_j. For a FIR filter, the filter output y_n is

$$y_n = \sum_{j=-\left(\frac{N-1}{2}\right)}^{\left(\frac{N-1}{2}\right)} h_j x_{n-j} \tag{4}$$

The output of a halfband filter is found by invoking (3), and y_n can be expressed as

$$y_n = \sum_{\substack{j=-\left(\frac{N-1}{2}\right) \\ j \neq 0,\, j \text{ odd}}}^{j=\left(\frac{N-1}{2}\right)} h_j x_{n-j} + x_n \tag{5}$$

In this expression for y_n with even n (decimation factor of 2), the only term in the output that involves the even samples is the last term, which corresponds to the center tap of the filter. Thus, it is possible to reduce the computation of y_n for even-n into a lower-order (half) filtering and an addition operation. First, the input is demultiplexed into odd and even samples. The odd samples are filtered using an $((N+1)/2)$-tap FIR filter. The filter output is summed with the delayed even sample. In Fig. 5(b), the simplicity of this architecture is evident when comparing the original FIR structure to its equivalent model. The length of the delay in the simplified model is

$(N + 1)/2 + \tau$, where τ is the pipeline delay through the FIR filter.

Hardware realization uses an $(N + 1)/2$ order FIR filter, a shift register for the delayed even samples, and an adder for adding the delayed samples to the filter output. This structure has the important additional advantage of operating the FIR filter at half the input frequency. Thus slower components (multiply accumulators) than a general FIR filter can be used to implement this architecture, which directly impacts the cost. Furthermore, it is also possible to take advantage of the symmetry of the FIR filter coefficients [7] to reduce the amount of computation by another factor of two.

B. Halfband Filter Design

Some examples of halfband filters are shown in Table 1. These were obtained using the Remez exchange algorithm [7]. In this article, we do not discuss the filter design package for this algorithm. Our design package uses the standard linear phase filter design package in [7], which can be used to design a variety of types of FIR filters.

In Table 1, the coefficients of three 15th order halfband filters are listed in the first 4 columns. Only four coefficients need to be specified, since $h_{-k} = h_k$, $h_0 = 1$, and $h_k = 0$ for k-even. The 5th column, f_p, is the upper band edge of the passband, expressed as a fraction of f_s. The stopband starts at $0.5-f_p$ and the transition band is from f_p to $0.5-f_p$. The last column is the deviation of the filter from ideal.

The frequency response of the 15th order filter corresponding to the first row of Table 1 is shown in Fig. 6. For the filter of Fig. 6, each coefficient is quantized to 9 bits of resolution, which is the coefficient accuracy of the FIR filter chip used in the DSN Advanced Receiver. With the resolution of the figure, there is no increase in the stopband deviation relative to the −38.6 dB predicted in Table 1. This filter results in a $0.6 * 10^{-3}$ dB loss due to aliasing and the same amount due to passband ripple loss [9]; thus, it results in a total loss of $1.2 * 10^{-3}$ dB.

The halfband filter was built and tested using the architecture depicted in Fig. 5(b). The FIR filter is implemented using an 8th order FIR filter chip, ZR-33891 (Zoran Corp.), operating up to 25 MHz. This chip supports 9-bit data and 9-bit coefficients, and it outputs 26 bits. The block diagram of the realization is shown in Fig. 7(a), and a photograph of the board is shown in Fig. 7(b).

For proper alignment of the 26-bit output of the FIR chip and the full adder, it is clear from Eq. (5) that the center tap corresponding to the delayed samples is "1," and therefore the output result of the FIR filter must be added starting at the 9th bit position.

C. VLSI Chip for Halfband Filters

A custom VLSI chip is under development to reduce the chip count in the halfband filter from 33 chips to 2 chips. The custom VLSI chip complements the ZORAN-33891 FIR chip. This chip operates in a fully synchronous mode from an input clock with frequencies of up to 25 MHz. The block diagram of the filter using this chip with the Zoran chip is shown in Fig. 8. It is noted that the demultiplexer at the input of the filter is not included in the VLSI chip. This is because of the demodulator design, wherein the odd and even samples are demultiplexed at the output into two paths to enable the use of slower (25 MHz) multipliers than would be required otherwise.

A detailed specification of this VLSI chip is contained in the Appendix.

IV. Dynamic Range and DC-Offset Analysis

A fundamental problem in the design of a digital mixer using finite precision arithmetic, especially fixed point arithmetic, is that a nonzero DC-offset (mean) is always present from quantization error, rounding, and truncation. The DC-offset is very undesirable in the process of complex heterodyning, since it directly translates into static phase error for residual carrier systems.

A. Maximum Allowable DC-Offset

In a residual carrier system, the values of the Q-channel samples at the demodulator output are used to estimate the carrier phase. When the signal-to-noise ratio is small, even a small DC-offset results in a large phase error. To explicitly exhibit this fact, denote the carrier power as P_c, the noise power as N_0, the noise variance as σ^2, the bandwidth as B, and the DC-offset as β. Then, for $P_c/N_0 = 0$ dB-Hz, the required threshold for the advanced receiver, and for $B = 15$ MHz, the carrier power signal to noise ratio is $P_c/\sigma^2 = P_c/(N_0 B) = -72$ dB. The signal amplitude is given by $A = \sqrt{2P_c}$. The total power is dominated by noise, and its scaling is such that the 8-bit A/D converter saturates at 4σ. Thus $\sigma = 2^5 = 32$, and A is 0.01, in units of the A/D output step size. With slowly varying random phase, the expected value of the complex mixer output is $A * \sin(\phi) + \beta$. Approximating $A \sin(\phi) + \beta$ with $A\phi + \beta$, the average phase estimate is $\hat{\phi} = 1/A(A\phi + \beta)$. Therefore, the static phase error due to DC-offset is roughly β/A radians. Hence, to keep the phase error under 0.1 radian, it is necessary to keep the dc offset under 0.001, or 90 dB below the total input power.

B. Quantization and Scaling

To analyze the effect of quantization on the performance of the digital receiver, we consider the quantization and scaling model shown in Fig. 9 and the A/D converter characteristic shown in Fig. 10.

1. Input quantization and scaling. The input signal $r(t)$ passes through an AGC amplifier which controls the total power. The A/D converter (sampler and quantizer) quantizes to 8 bits, covering the range -2^7 to $2^7 - 1$ in quantizer output units. The nominal scaling is such that the total power at the quantizer output is 2^{10}, i.e., saturation is at approximately 4σ.

The FIR filter chip can accept only 9-bit inputs, which motivates the scaling at the multiplier output. We discuss the scaling to obtain 9-bit outputs, although scaling down to 8 bits is also evaluated.

The A/D output is multiplied by the in-phase and quadrature reference signals, which are stored in read only memory (ROM). These are stored as k-bit integers, $k = 12$ or 16. In either case, the amplitude is 2^{15}, i.e., the k most significant bits of the 16-bit multiplier input are used. The input signal from the A/D converter is scaled up by four at the input to the multiplier. This scaling results in a maximum possible value at the multiplier output of 2^{17} times the signal input value, i.e., a total maximum value of less than 2^{24}.

The multiplier chip has a feature which allows rounding from the least significant 16 bits into the next bit. This rounding feature is used, effectively scaling the level down by 2^{16}. Thus the final output of the multiplier, after rounding, is in the range -2^8 to $2^8 - 1$, and is represented by $L = 9$ bits. This forms the input to the FIR filter. It is the DC-offset at this point in the system which dominates static phase error performance.

2. Signal scaling. The net effect of the above on signal level is to multiply the input Q-channel signal by $2 \sin(\omega_c n T_s + \hat{\theta})$. Expanding the sum and difference frequency terms of the multiplier output, the factor of 2 is cancelled. The difference frequency term at the multiplier output has the same signal amplitude and the same noise power as the input process, i.e., the Q-channel signal is $A \sin(\hat{\theta} - \theta)$ and the noise power in the low frequency process is σ^2.

The approximate effect of the FIR filter is to pass all of the low frequency terms, amplified by the DC gain of the filter. The DC gain of the filter, in amplitude, is the sum of the coefficients, i.e.,

$$\sum h_i$$

These coefficients can be expressed as 9-bit numbers, so the scaling is such that the center tap gain is 2^8, and the sum of h_i is approximately 2^9. Therefore, at the output of the FIR filter, the signal is

$$2^9 A \sin(\hat{\theta} - \theta)$$

where A is the signal value at the A/D converter output, and the noise power is $2^{18}\sigma^2$. The input scaling has been chosen to saturate at 4σ with an 8-bit A/D converter (at low SNR). Preserving the same saturation level at the output clearly requires 17 bits. The reason we also consider scaling the FIR filter input down by a factor of 2, to 8 bits, is so that only 16 bits would be required at the final output. Another way to get a 16-bit output would be to round or truncate the least significant bit of the 17-bit output, but this would introduce DC offset.

C. Sine and Cosine Table Accuracy

Demodulation is accomplished by multiplying the A/D converter output samples by cosine and sine tables. These components are prestored in read only memory (ROM) and are quantized to k bits. We specifically consider the case when $k = 12$ or 16. Besides having enough resolution in the sine and cosine tables, it is important that the N-point tables be designed such that the sum over the N-points is very close to zero. To satisfy this, consider the identity

$$\sum_{n=0}^{N-1} e^{j\omega_c n T_s} = \frac{1 - e^{j\omega_c N T_s}}{1 - e^{j\omega_c T_s}} \tag{6}$$

For this sum to be zero it is necessary that $\omega_c N T_s \cong 2\pi n$, for some integer n. However, even if N satisfies this requirement, due to quantization (rounding in this case) to k-bits of each value in the ROM, this sum is never exactly zero.

D. Evaluation of DC-Offset

To evaluate the DC-offset, we compute the sample mean at the output of the mixer. The received signal is the sum of the transmitted signal plus Gaussian noise. If we were to compute the sample mean by using Monte Carlo type simulation, the variance of the estimate would be $\sigma^2_{\text{out}} = \sigma^2/N$, where N is the number of independent noise samples. To get a sufficiently accurate estimate we need $\sigma_{\text{out}} \ll 10^{-3}$, the required DC-offset. For $\sigma_{\text{out}} = 10^{-4}$, it would be necessary to process $N = \sigma^2/\sigma^2_{\text{out}} = 2^{+10}/10^{-8} \cong 10^{11}$ samples. This is clearly too many samples to process on a general purpose computer.

However, it is possible to analytically derive the expression for the probability distribution function defined over the dynamic range of every stage of our system, since the number

of quantization intervals is small in our case. The mean and the variance of the signal are then computed from the probability distribution.

1. Probability distribution of the mixer output. In order to analyze the distribution of the dynamic range of the demodulated signal, we consider the received signal $r(t)$ in Fig. 9. Define $\Omega = \{-M, -M+1, \ldots, 0, 1, \ldots, M-1\}$ where $2 * M$ denotes the number of quantization intervals at the output of A/D. The probability space is (Ω, F, P), where F is the collection of all subsets of Ω. In this section, we drive the first order discrete probability density function for $y_n = T(y(nT_s))$, where $T(\cdot)$ denotes the quantizer characteristic function shown in Fig. 10. Here $\{y_n\}$ is a stationary discrete-time and discrete-valued stochastic process, generated by sampling and quantizing the continuous-time and continuous-valued stochastic process $y(t)$ for $t > 0$; i.e., the discrete-time process $y(nT_s) \in \mathcal{R}_1$ (prior to quantization) is a Gaussian process in this case [2], [3], with mean m and variance σ^2.

Let P_k denote the probability of the kth quantization interval, i.e., $P_k = Pr(y(nT_s) \in A_k)$, for $k = -M, -M+1, \ldots, 0, 1, \ldots, M-1$, where $2 * M$ is the number of quantization levels. Let

$$A_k = [k\Delta, (k+1)\Delta), \text{ for } -M < k < M-1$$

$$A_{-M} = [-\infty, -M\Delta), \text{ and } A_{M-1} = [M\Delta, +\infty)$$

Notice $\mathcal{R}_1 = \bigcup_{k=-M}^{M} A_k$. The probability of the received signal falling in the kth quantization interval is

$$P_k = Pr(y(t) \in A_k) = \int_{A_k} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(u-m)^2}{2\sigma^2}} du \qquad (7)$$

Let

$$erfc(x) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt$$

Then

$$P_k = erfc\left(\frac{(k+1)\Delta - m}{\sigma}\right) - erfc\left(\frac{k\Delta - m}{\sigma}\right) \qquad (8)$$

The probability distribution of y_i is then $Pr(y_i = k) = P_k$. Figure 11(a) and (b) shows the input probability density and the cumulative probability distribution when the signal has zero mean and unit variance (= 2^{10} for 8-bit input data).

Next, we compute the probability distribution at the output of the multiplier, which multiplies the A/D output y_n by the value s_n, the output of the sine (or cosine) table, and then quantizes the product. The output is denoted by $w_n = Q(s_n y_n)$, where $Q(\cdot)$ denotes the quantization function, at the multiplier output. The quantizer $Q(\cdot)$ for fixed point arithmetic is truncation or rounding. The function $Q(\cdot)$ depends on how the output product is rounded. In our hardware implementation a standard 16-bit multiplier is used, which accepts two 16-bit inputs in two's complement and outputs a 31-bit product. The product can be rounded only at one bit position, namely from the 15th bit into the 16th bit, where bits are labeled starting at zero for the least significant bit. We scale the two input numbers such that $L = 8$ or 9 bits remain, because the filter which follows the multiplier can accept only 9 bits of input.

The sequence s_n is quantized to k-bits. Thus let

$$s_n \in \Omega' = \{-2^k + 1, \ldots, 2^k - 1\}$$

and

$$w_n \in \Omega'' = \{-2^L + 1, \ldots, 2^L - 1\}$$

Then $Pr(w_n = \alpha)$ may be expressed as

$$Pr(w_n = \alpha) = \sum_{\xi} Pr(Q(\xi y_i) = \alpha) Pr(S_n = \xi) \qquad (9)$$

since s_n's and y_n's are independent. The sequence s_n is selected uniformly from an alphabet of size N, with k-bits of resolution; thus, $Pr(s_n = \xi) = 1/N$, and $Pr(Q(\xi y_n) = \alpha)$ is given by (8), for each $y_n = k$ and $\xi \in \Omega'$.

Figure 12 shows the output densities for three cases. Figure 12(a) and (b) is for the in-phase and quadrature components, respectively, for $L = 8$ and $k = 12$. Figure 12(c) is for the in-phase component, when $L = 8$ and $k = 16$. Figure 13 shows the input and output densities for a case of improper input scaling. In particular, the input signal level is higher by a factor of 4 in power, i.e., the input power is 2^{12}. The two spikes of the input probability distribution function are due to the improper scaling and consequent saturation.

2. Results for computation of DC-offset. In Table 2, the values of the induced DC-offset due to sine/cosine quantization (for $N = 99$, which is the number of points in the ROM table, and $\omega_c T_s = 1/3.96$) and multiplier output quantization

are given. This table lists the results for both the 12- and 16-bit sine and cosine tables, and with 9- and 8-bit quantized output. These values were computed by calculating the mean of the probability distribution functions given in Eqs. (8) and (9). Values in the table are in units at the output of the multiplier, scaled between 1 and −1, i.e., scaled down by 2^8 from the integer representation.

The first column is the resolution of the ROM table. The second column is the number of bits at the output of the quantizer. The last two columns give the resulting DC-offset at the output of the multiplier, corresponding to each arm. This corresponds to the DC-offset bias in static phase error, for the sine table, and to error in measuring A, for the cosine table. The gain of the system is such that the signal amplitude is the same at the multiplier output as the A/D converter output, i.e., it is 0.01 integer units for P_c/N_0 = 0 dB-Hz. Scaling the multiplier output to the range −1 to +1 reduces the signal amplitude by a factor of 2^8, to approximately $4 * 10^{-4}$. Referring to Table 2, the static phase error is small, i.e., less than $4 * 10^{-5}$, for $k = 12$, $L = 8$ and for $k = 16$, $L = 9$, and is marginal for $k = 16$, $L = 8$.

The DC-offset for the in-phase arm is not acceptable for $k = 12$, but is good for $k = 16$. For this reason we have elected to use the 16-bit resolution table, with 9-bit multiplier output resolution.

3. Effects of FIR filter on static phase error. In our hardware design each product $h_{i-n} w_i$ is computed and accumulated without any quantization. Therefore, the mean value at the output of the filter is exactly equal to $\sum_{i=0}^{N-1} h_i$ times the mean value at the input. The input mean is $A\phi + \beta$, where β is the DC-offset and the static phase error is β/A. Since both β and A are multiplied by the same filter gain, there is no change in static phase error in the filter. Some additional static phase error is introduced if the filter output is quantized in further processing.

V. Conclusions

In this article we outlined the design and analysis of a digital complex demodulator to be used in applications that require high speed (up to a 50-MHz sample rate) and operate at very low SNR. We found a cost-effective way to implement the digital low-pass filter required for this application. A two chip solution with the architecture of a custom VLSI coprocessor chip was proposed. The filter results in less than $1.2 * 10^{-3}$ dB loss due to passband ripple and imperfect attenuation of the unwanted sum frequency terms at the output. A comprehensive analysis of the dynamic range distribution was presented, using a probabilistic approach, and it was deduced that it is possible to attain a static phase error of less than 0.2 deg at the DSN advanced receiver threshold SNR of P_c/N_0 = 0 dB, with a 15-MHz bandwidth.

References

[1] J. H. Yuen (ed.), *Deep Space Telecommunications Systems Engineering*, JPL Publication 82-76, Jet Propulsion Laboratory, Pasadena, California, July 1982.

[2] W. C. Lindsey and M. K. Simon, *Telecommunications Systems Engineering*, Englewood Cliffs, New Jersey: Prentice-Hall, 1973.

[3] R. E. Zeimer and R. L. Peterson, *Digital Communications and Spread Spectrum Systems*, New York: Macmillan Publishing Co., 1985.

[4] R. Sfeir, S. Aguirre, and W. J. Hurd, "Coherent Digital Demodulation of a Residual Carrier Signal Using IF Sampling," *TDA Progress Report 42-78*, vol. April–June 1984, Jet Propulsion Laboratory, Pasadena, California, pp. 135–142, August 15, 1984.

[5] R. Sadr and W. J. Hurd, "Detection of Signals by the Digital Integrate-and-Dump Filter With Offset Sampling," *TDA Progress Report 42-91*, vol. July–September 1987, Jet Propulsion Laboratory, Pasadena, California, pp. 158–173, November 15, 1987.

[6] R. Sadr, "Detection of Signals by Weighted Integrate-and-Dump Filter," *TDA Progress Report 42-91*, vol. July–September 1987, Jet Propulsion Laboratory, Pasadena, California, pp. 174–185, November 15, 1987.

[7] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*, Englewood Cliffs, New Jersey: Prentice-Hall, 1975.

[8] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*, Englewood Cliffs, New Jersey: Prentice-Hall, 1983.

[9] R. Sadr and W. J. Hurd, "Filter Distortion Effects on Telemetry Signal-to-Noise Ratio," *TDA Progress Report 42-88*, vol. October–December 1986, Jet Propulsion Laboratory, Pasadena, California, pp. 59–66, February 15, 1987.

[10] M. G. Bellanger and G. Bonnerot, "Interpolation, Extrapolation, and Reduction of Computation Speed in Digital Filters," *IEEE Trans. ASSP*, vol. 25, no. 4, pp. 231–235, August 1974.

[11] M. G. Bellanger, "Computation Rate and Storage Estimation in Multirate Digital Filtering with Half-Band Filters," *IEEE Trans. ASSP*, vol. 25, no. 4, pp. 344–346, August 1977.

[12] D. H. Brown and W. J. Hurd, "DSN Advanced Receiver: Breadboard Description and Test Results," *TDA Progress Report 42-89*, vol. January–March 1987, Jet Propulsion Laboratory, Pasadena, California, pp. 48–66, May 15, 1987.

Table 1. Examples of halfband filters with $N = 15$

Coefficients				Passband	DEV-dB
h_1	h_3	h_5	h_7	f_p	
0.312	−0.089	0.038	−0.017	0.187	−38.62
0.309	−0.081	0.029	−0.085	0.162	−51.25
0.312	−0.078	0.026	−0.006	0.150	−57.93

Table 2. Mean values (all numbers scaled to range ±1)

ROM resolution, k	Output resolution, L	Average, sine table	Average, cosine table	DC in-phase arm	DC quadrature arm
12	8	2.42×10^{-4}	2.86×10^{-4}	4.71×10^{-4}	5.48×10^{-6}
16	8	1.57×10^{-5}	1.36×10^{-5}	-2.98×10^{-6}	-6.15×10^{-5}
16	9	1.57×10^{-5}	1.36×10^{-5}	-6.55×10^{-6}	-3.57×10^{-5}

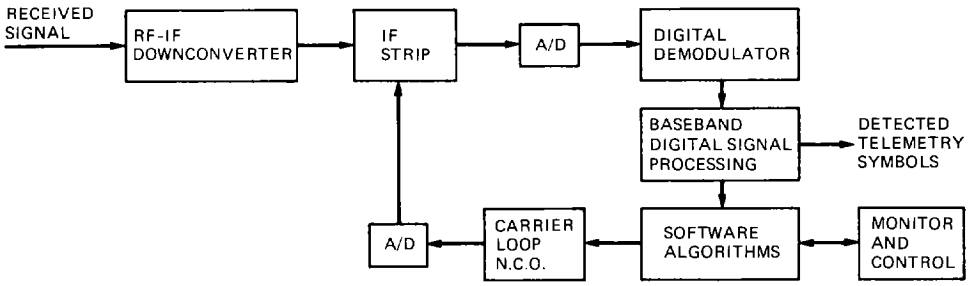

Fig. 1. Simplified architecture of advanced receiver

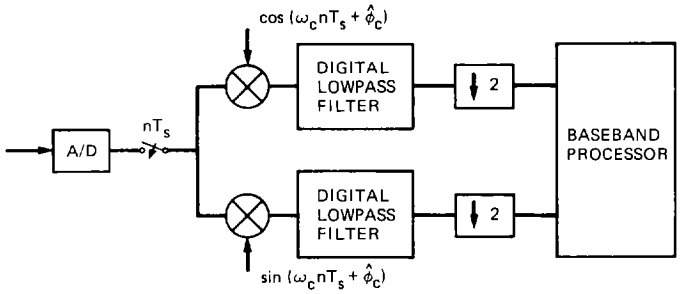

Fig. 2. Complex demodulator

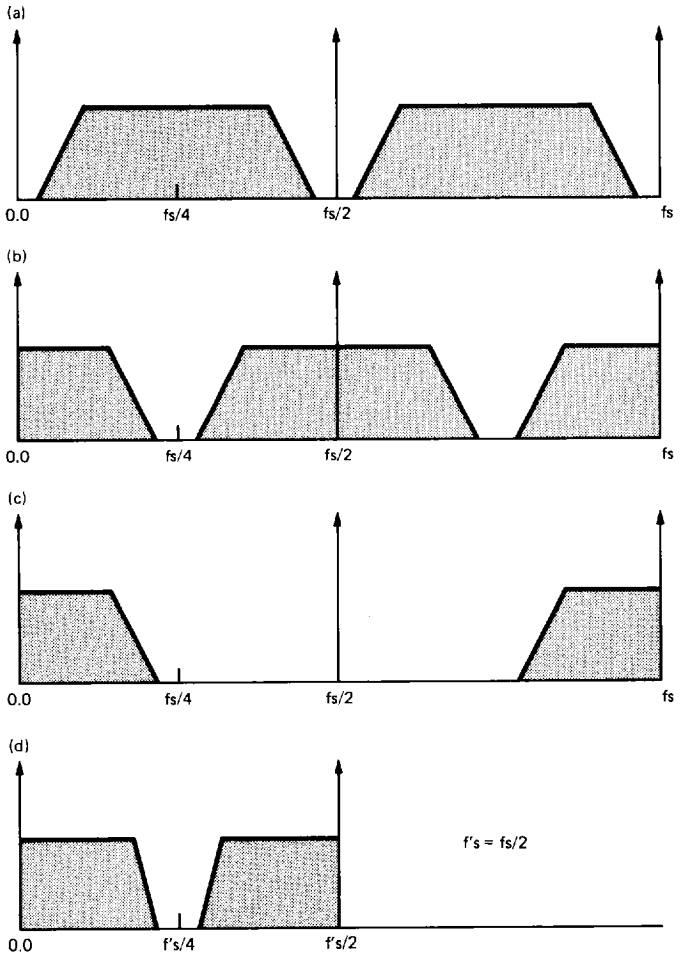

Fig. 3. Spectrum of signals: (a) A/D output; (b) mixer output; (c) filter output; and (d) decimated output

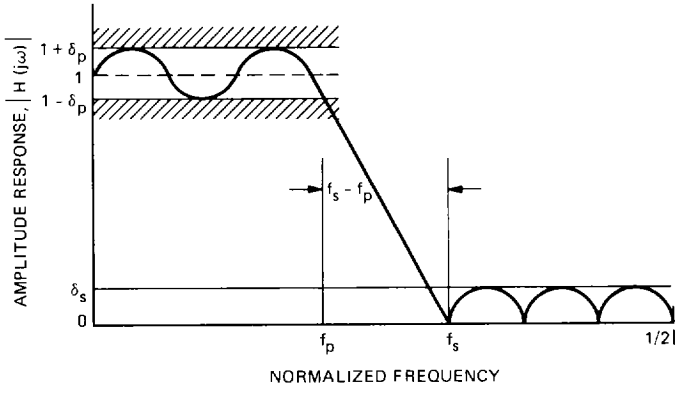

Fig. 4. Tolerance parameters for a practical FIR lowpass filter

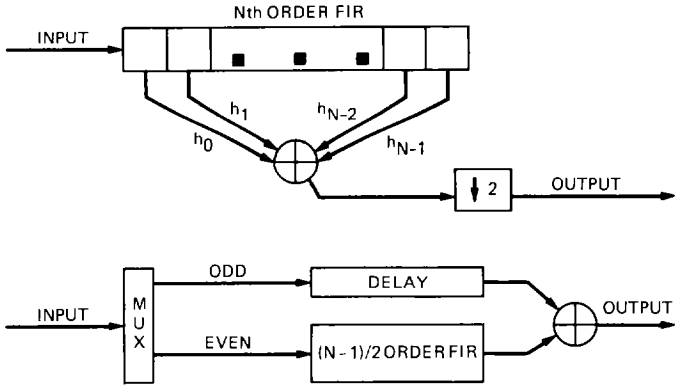

Fig. 5. FIR: (a) *N*th order FIR with decimation factor 2; (b) equivalent halfband model

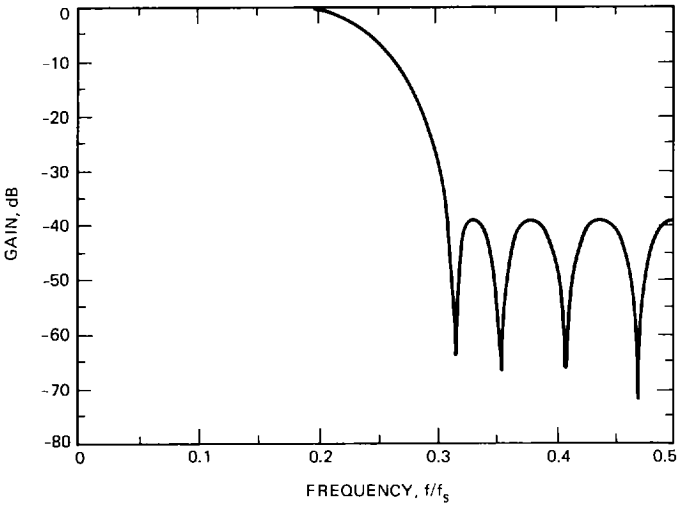

Fig. 6. Frequency response of a 15th order halfband filter

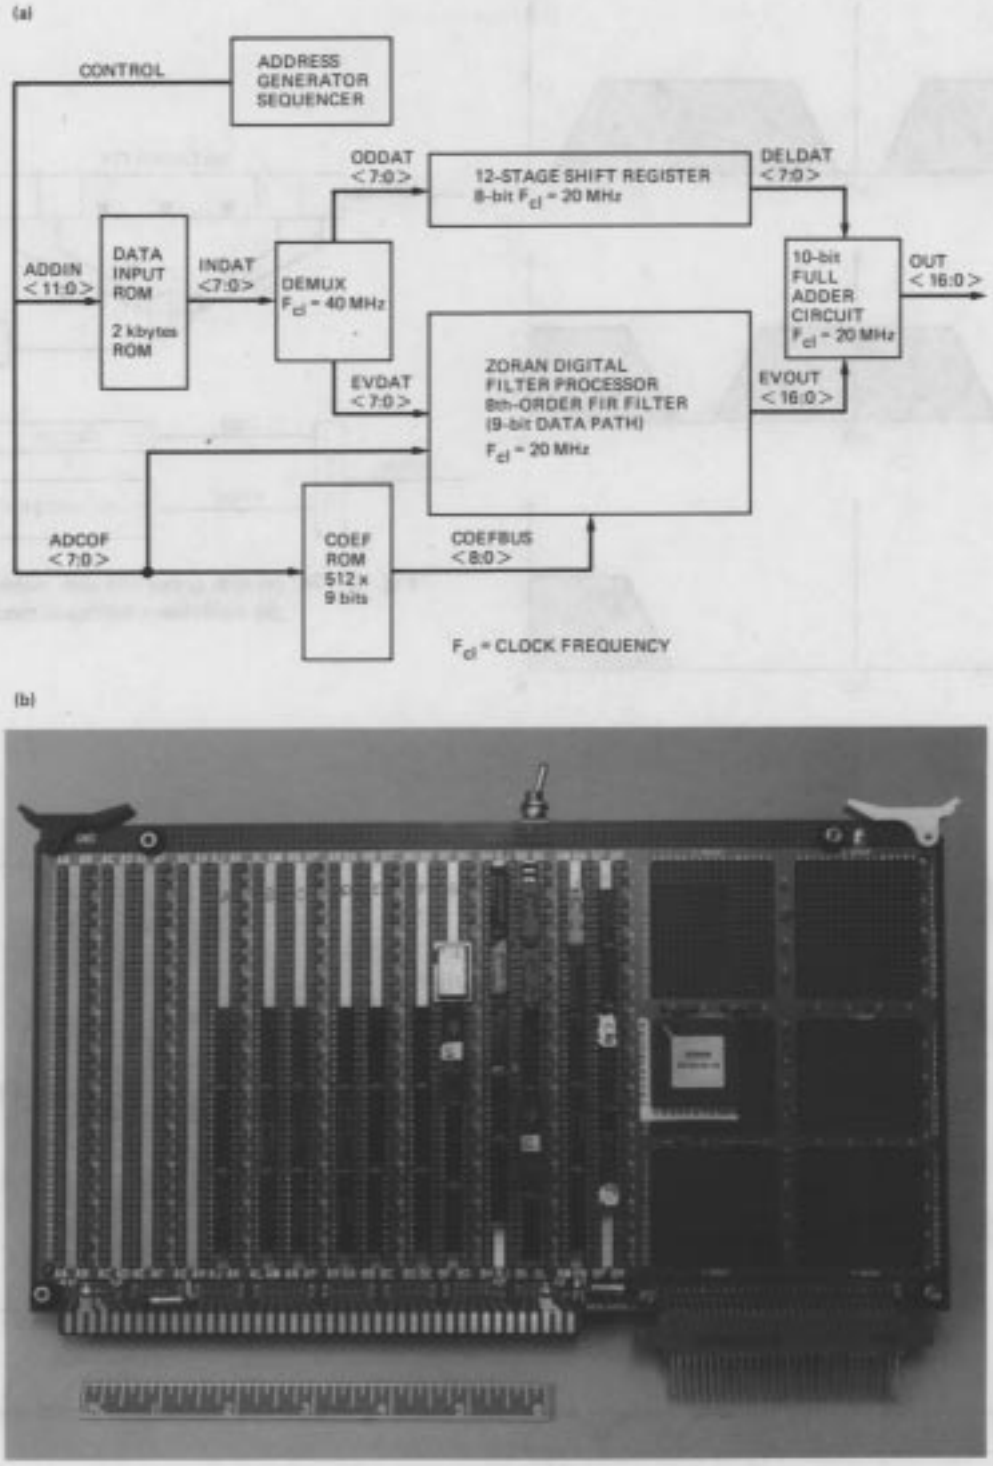

Fig. 7. Hardware realization of halfband filter: (a) block diagram; (b) breadboard

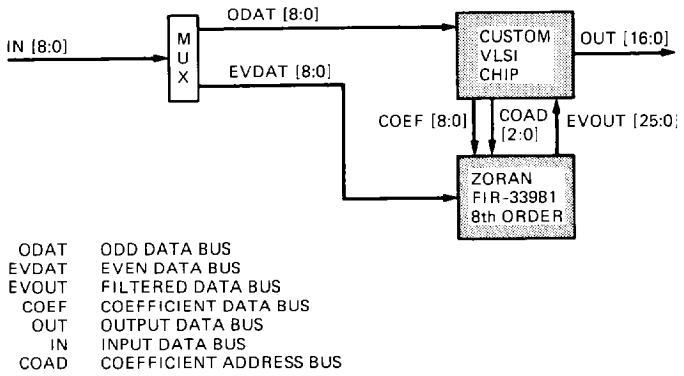

Fig. 8. Custom VLSI coprocessor and ZORAN chip

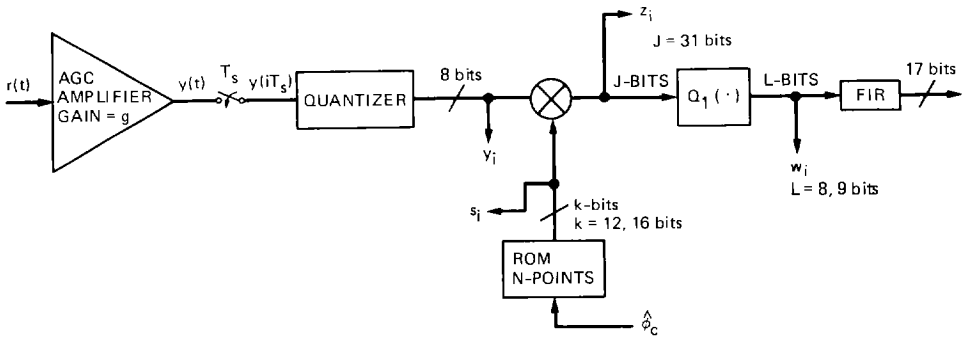

Fig. 9. Quantization model for analysis

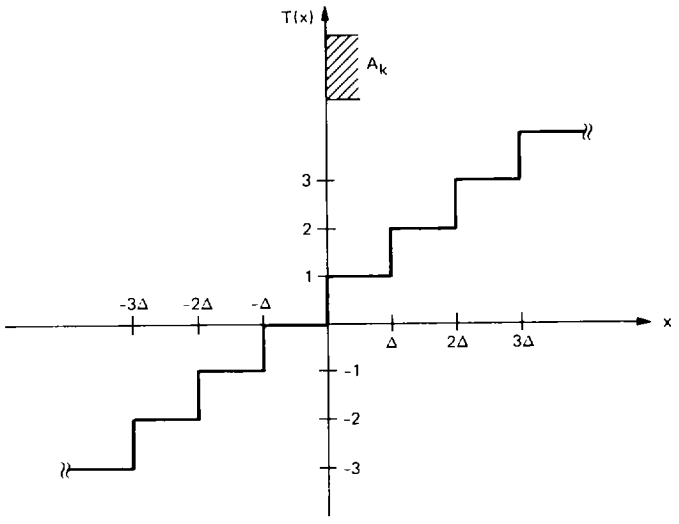

Fig. 10. A/D converter characteristic

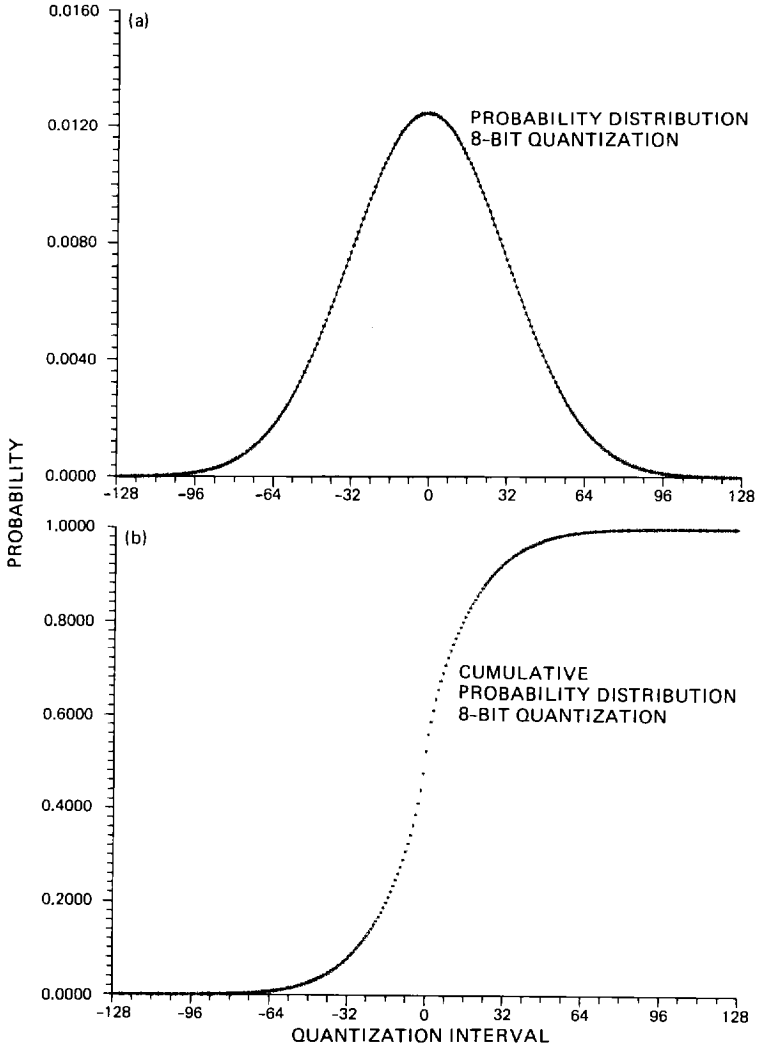

Fig. 11. Probability functions for A/D converter output

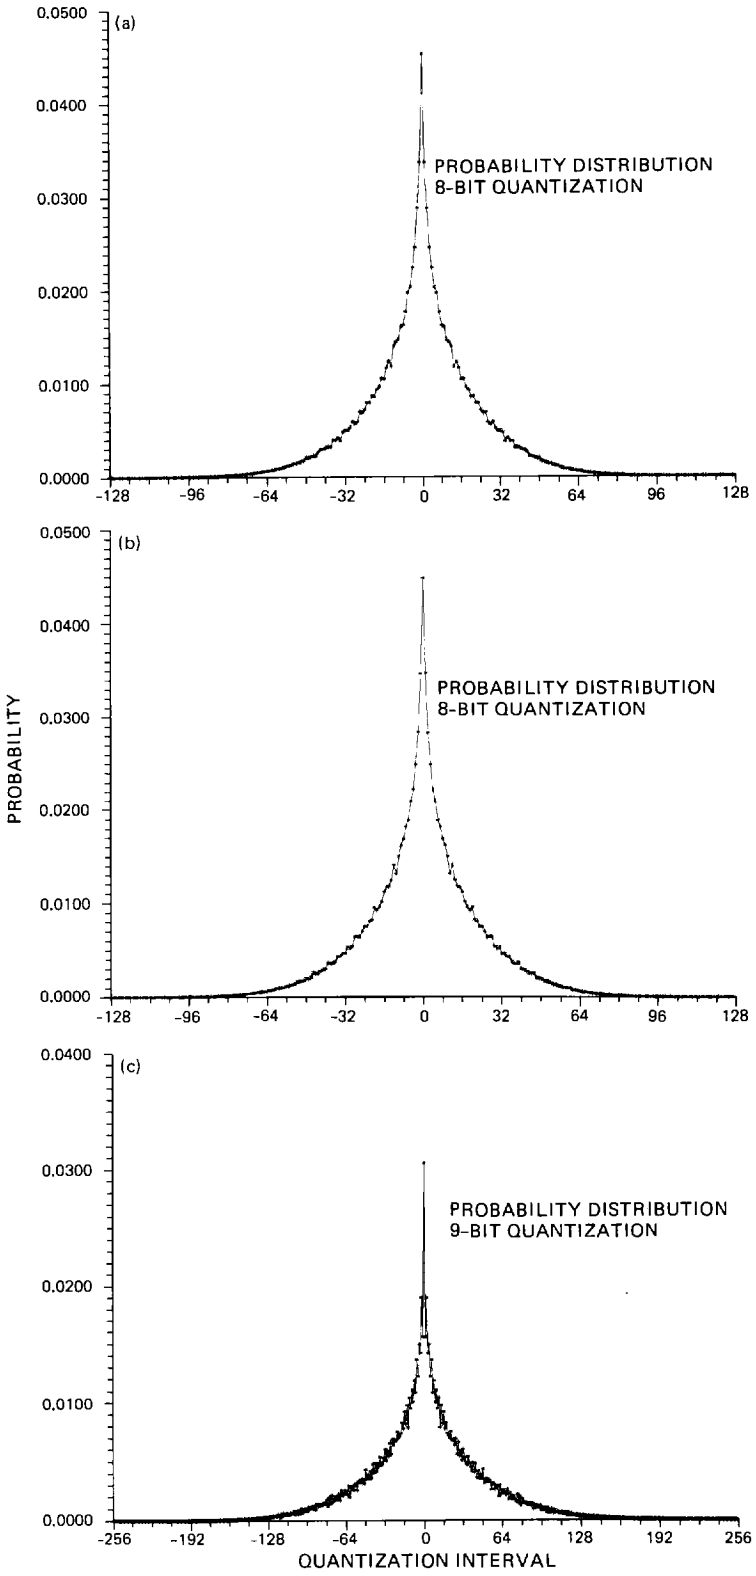

Fig. 12. Probability densities at multiplier output for normal scaling: (a) in-phase arm; (b) quadrature arm (k = 12, L = 8); and (c) in-phase arm (k = 16, L = 9)

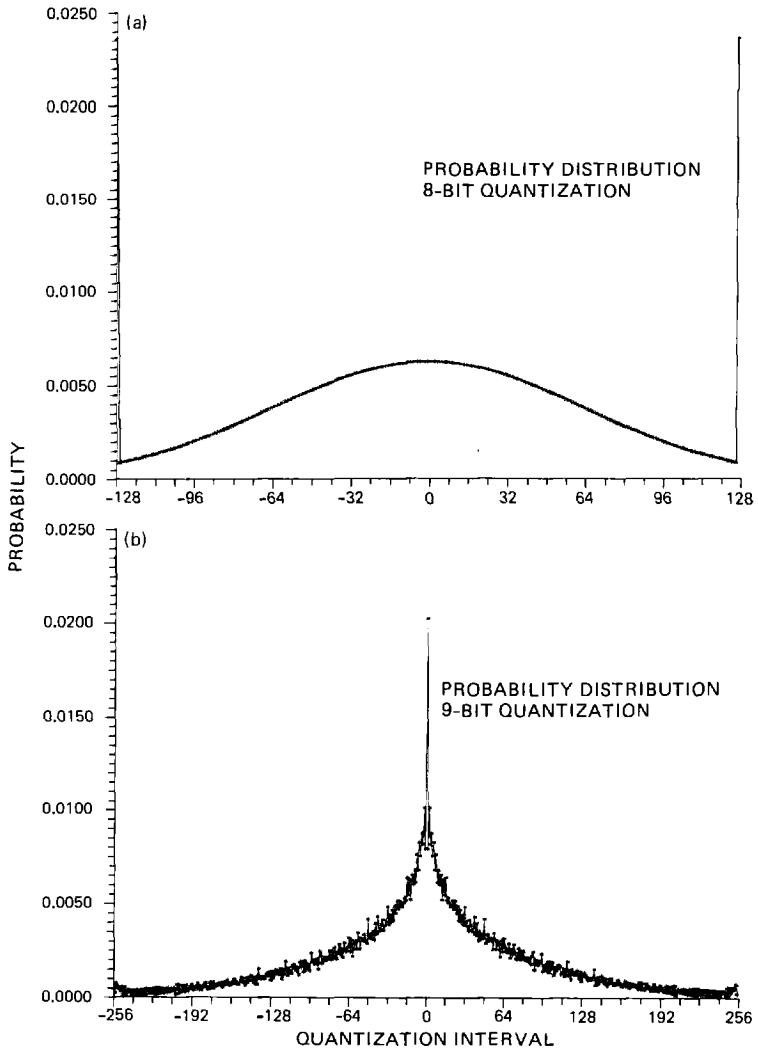

Fig. 13. Probability densities at the A/D converter for signal scaled for input power of four times normal: (a) input density (g = 1/2); (b) output density (k = 16, L = 9, g = 1/2)

Appendix
Chip Specification

In this appendix we outline the top-level specification of the design of a VLSI chip that complements the FIR chip, for a two chip solution of the halfband filter discussed in Section III. The functional block diagram of this chip is shown in Fig. A-1(a).

This chip operates in a fully synchronous mode from an input clock with frequencies up to 25 MHz. The transfer of data from the input, through all subsequent stages within the pipeline, and to the output stage is performed on the rising edge of the input clock. The 9-bit input data is delayed by a 12-stage shift register. The delayed data ODATD is added to the input from the ZORAN chip. As shown in Fig. A-1(b), this chip contains an 8- by 9-bit-wide RAM that can be preloaded by the host. In the normal mode of operation, 8 locations of the RAM are sequentially accessed, providing the coefficient stream to the ZORAN part.

For host communication, during the initialization period, the input address strobe signal /AS (/ stands for active low) is asserted by the host to initiate a bus write cycle to the RAM. The output data transfer acknowledge signal /DTACK, which is generated on-chip, is asserted when the input data is latched, so that the host may terminate the write cycle and thus negate /AS. These two signals, and the input clock, are sufficient to provide both synchronous or asynchronous write cycles to the chip, depending on the type of host processor used to interface this chip.

The 26-bit filtered data, EVOUT, is aligned with the 9-bit ODDAT as shown in Fig. A-1(b). The least significant 7 bits are simply delayed by a pipeline delay to compensate for the delay introduced by the output stage. Starting from the 9th bit (EVOUT-8), 9 bits are aligned to the input of the full adder, and the remaining bits are discarded because they are just sign extensions.

The output of the 10-bit full adder, shown in Fig. 11(b), is overflow protected at 17 bits, to prevent wrap-around when using the 2's complement number presentation. A 17-bit output results.

We have listed the pin list and the description of each signal. Total pin count of this chip is 78. The timing specification of the signals and a detailed discussion of the operation of this chip are deferred to a later article.

Table A-1. VLSI input/output definitions

Signal	Width	I/O	Name
CS	1	I	Chip Select
ODAT	9	I	Odd Data Bus
EVOUT	17	I	Filtered Even Data Bus
OUT	17	O	Output Data Bus
INH	9	I	Host Data Bus
/AS	1	I	Host Address Strobe
/DTACK	1	O	Data Transfer Acknowledge
INHA	3	I	Host Address Bus
COEF	9	O	Coefficient Data Bus
CLK	1	I	Input Clock
RESET	1	I	Reset Signal
VSS	4	X	+5 Volt Power Supply
GND	5	X	Ground Terminal

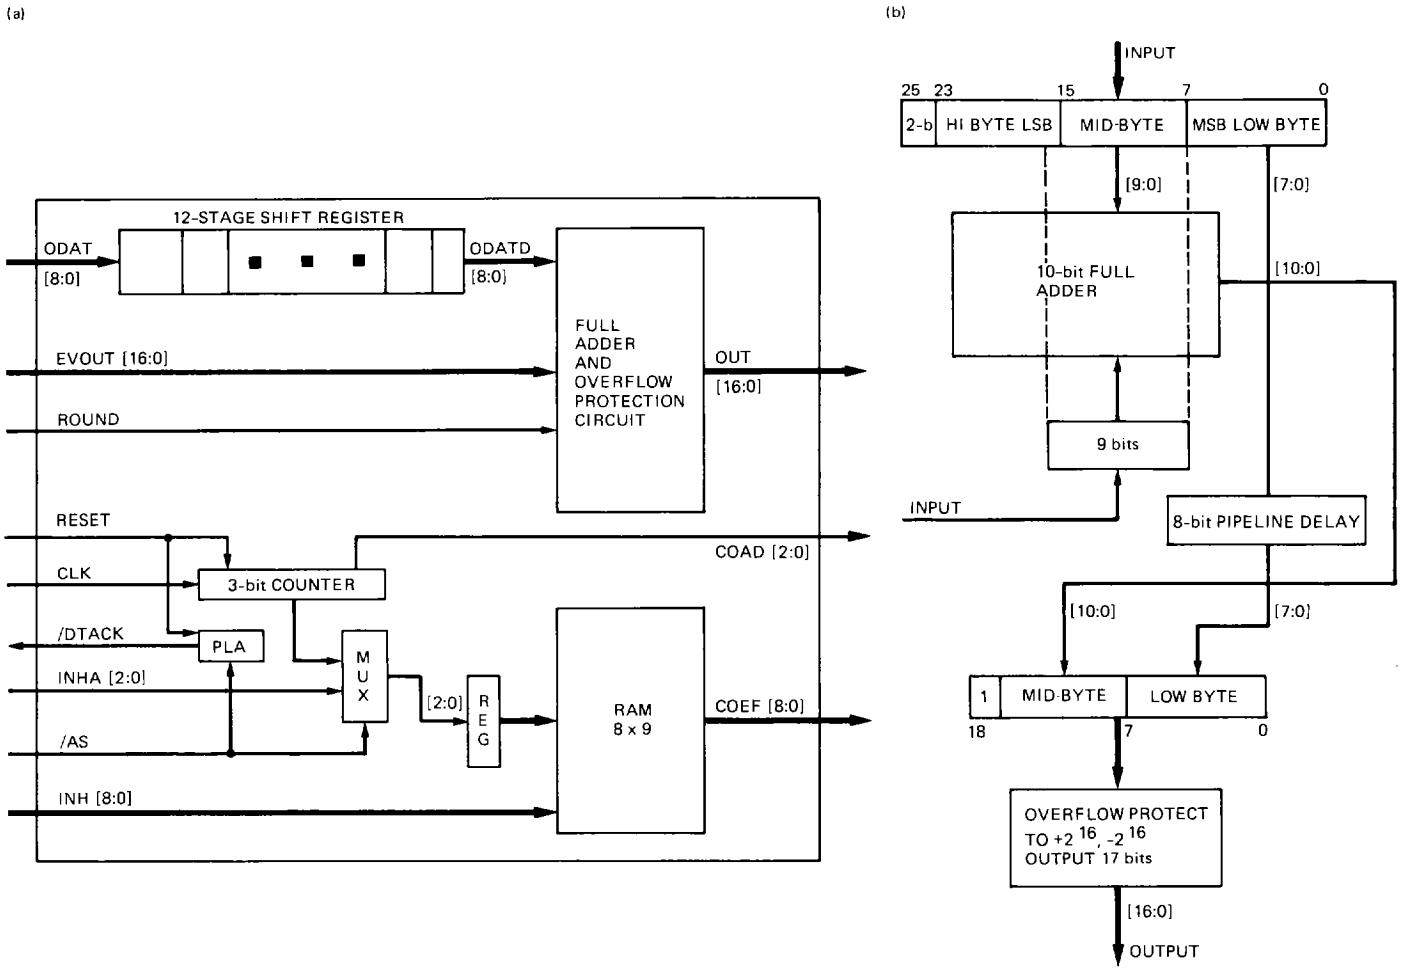

Fig. A-1. Custom VLSI chip: (a) functional block diagram; (b) full adder