Integrated Circuits for Data Transmission Over Twisted-Pair Channels

David A. Johns and Daniel Essig

***Abstract*— This tutorial paper discusses typical architectures and challenges in designing integrated circuits for data transmission over twisted-pair wire channels. To highlight the various architectural approaches, two main applications are discussed—high-bit-rate digital subscriber loop (HDSL) and fast-ethernet. Although these two applications have orders of magnitude difference in their bit rates, they share many common building blocks including line-drivers, 2–4 wire hybrids, echo cancellation, digital equalization, and clock recovery. Typical integrated circuit approaches for realizing each of these blocks are presented as well as possible tradeoffs. Finally, future challenges facing integrated circuit designers are presented.**

***Index Terms*— Clock recovery, digital transmission, echo cancellation, equalization.**

I. Introduction

DUE to the abundance and low-cost of unshielded twisted-pair (UTP) cables, there is great interest in transmitting high-speed data over UTP cable. However, there are certain challenges that face circuit and system designers in accomplishing this task. This tutorial paper intends to highlight the various architectural approaches used for overcoming some of the challenging design obstacles.

While common circuit and architectural approaches can be used in a variety of applications, for illustrative purposes, it is useful to focus on specific applications. Specifically, two main applications are addressed here—high-bit-rate digital subscriber loop (HDSL) and fast-ethernet. Other twisted-pair data-communication applications of interest but not specifically addressed in this paper include ISDN, asymmetric digital subscriber loop (ADSL), E1/T1, and emerging ATM standards.

HDSL and fast-ethernet make use of a variety of main blocks as shown in Fig. 1. The transmit D/A is used to convert digital levels to a suitable analog signal. A low-impedance line-driver is used to supply the necessary drive currents to the cable. The 2–4 wire hybrid allows both transmission and reception of data over a single cable. The receive A/D converts the received signal (plus some unwanted transmitted signal) to a digital signal. The echo canceler is used to eliminate residual transmit signal on the receive path due to a nonideal 2–4 wire hybrid (as we shall see, practical hybrids only partially cancel the transmit echo). Finally, clock-recovery and equalization are performed on the received signal to determine the correct timing information and remove intersymbol interference (ISI).

Manuscript received August 1, 1996; revised October 30, 1996.

D. A. Johns is with the Department of Electrical and Computer Engineering, University of Toronto, ON, M5S 1A4, Canada.

D. Essig is with the Brooktree Corp., San Diego, CA 92121 USA.

Publisher Item Identifier S 0018-9200(97)01067-6.

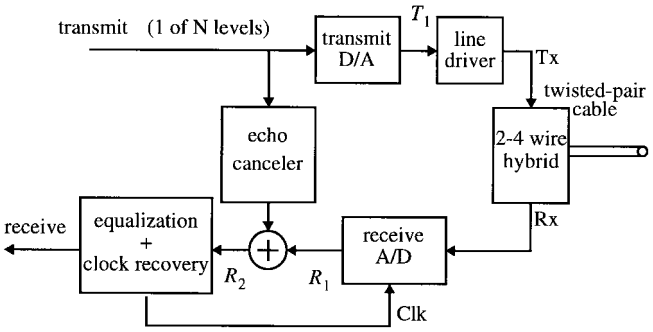

Fig. 1. The main blocks for twisted-pair data transmission.

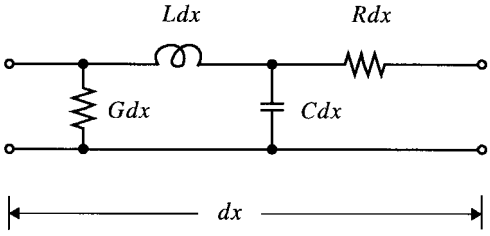

Fig. 2. A lumped-parameter model for a short section of cable.

The organization of this paper is as follows. First, modeling of twisted-pair cables is discussed. Next, the two main applications of HDSL and fast-ethernet are described. Each of the blocks shown in Fig. 1 are then discussed in more detail. Finally, future challenges of some of these blocks are described.

II. Twisted-Pair Cable Modeling

A twisted-pair cable is modeled as a transmission-line and can be described using four primary constants—R, G, L, C [1], [2]. These four "constants" are shown in Fig. 2 where (in terms of "per unit length"), R is the internal resistance, G is the conductance, L is the inductance, and C is the capacitance. Due to the "skin-effect" for ac signals, the internal resistance, R, is actually a complex impedance and can be modeled by [3]

$$R(\omega) = k_R\,(1+j)\,\sqrt{\omega}\,\Omega/\text{km} \tag{1}$$

where k_R is a constant determined by the diameter and material of the wires (ω is in units of rad/s). Note that $R(\omega)$ is proportional to the square root of frequency. Parameters L and C are relatively constant at higher frequencies and $G \cong 0$ in modern cables.[1] Typical values for these parameters

[1]The conductance does not equal zero in poorer quality cables, such as CAT3, which leads to a term proportional to f in the cable transfer-function.

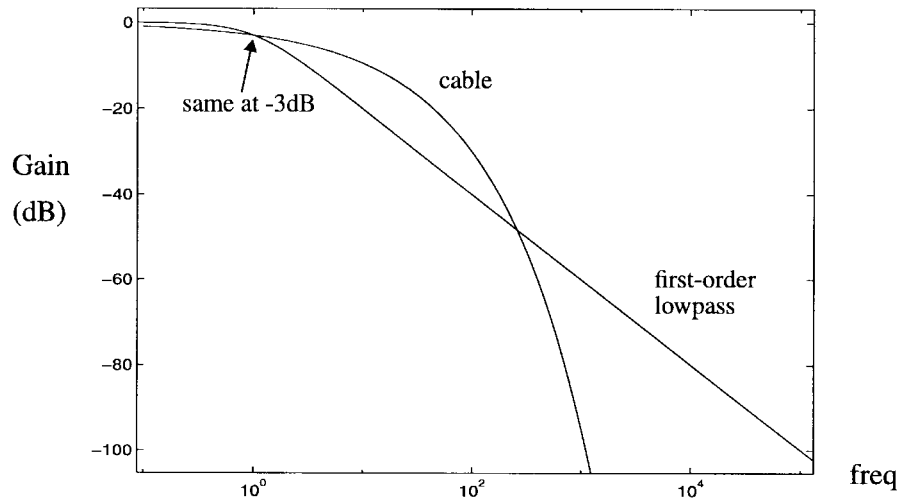

Fig. 3. Comparison of gain for an ideal cable and a first-order lowpass filter.

are $k_R = 0.2$, $L = 0.6$ mH/km, and $C = 0.05\,\mu$F/km at frequencies higher than 100 kHz.

A. Characteristic Impedance

The characteristic impedance of a transmission-line can be shown to equal

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}}. \tag{2}$$

Substituting in (1) $G = 0$ and using the approximation $\sqrt{1+x} \cong 1 + x/2$ for $x \ll 1$, we find

$$Z_0 \cong \sqrt{\frac{L}{C}} + \frac{k_R(1-j)}{2\sqrt{\omega LC}}. \tag{3}$$

Thus, for high frequencies, we have $Z_{0h} \cong \sqrt{L/C}$. In other words, since L and C are relatively constant, the characteristic impedance of the line is relatively constant at higher frequencies. Using the typical values above, we have $Z_{0h} = 110\,\Omega$ for frequencies greater than 100 kHz implying that when terminating the cable, impedances around 110 Ω should be used. At lower frequencies, the characteristic impedance has an extra term which is inversely proportional to $\sqrt{f}$.

B. Transfer-Function

When properly terminated, the transfer-function of a twisted-pair cable is modeled by

$$H(d, \omega) = e^{-d\gamma(\omega)} = e^{-d\alpha(\omega)}\, e^{-jd\beta(\omega)} \tag{4}$$

where $\gamma(\omega) = \alpha(\omega) + j\beta(\omega)$ and is given by

$$\gamma(\omega) = \sqrt{(R + j\omega L)\,(G + j\omega C)}. \tag{5}$$

Here, α and β are the attenuation and phase constants, respectively, while d is the cable length. Setting $G = 0$ and substituting (1) into (5), we can write

$$\gamma(\omega) = j\omega\sqrt{LC}\left[1 + \frac{k_R\,(1-j)}{L\sqrt{\omega}}\right]^{1/2}. \tag{6}$$

Now using the approximation $\sqrt{1+x} \cong 1 + x/2$ for $x \ll 1$ and making use of $\gamma(\omega) = \alpha(\omega) + j\beta(\omega)$, the following results are obtained:

$$\alpha(\omega) = \frac{k_R}{2}\sqrt{\frac{\omega C}{L}} \tag{7}$$

$$\beta(\omega) = \omega\sqrt{LC} + \frac{k_R}{2}\sqrt{\frac{\omega C}{L}}. \tag{8}$$

Combining (4) and (7), we see that for a given length d, the cable's transfer-function gain (in dB) is given by

$$\begin{aligned} H_{\mathrm{dB}}(d, f) &= 20\log_{10}|H(d, f)| \\ &= \frac{-20}{\ln 10}\, d\alpha(f) \\ &= -8.686d \times k_R\sqrt{\frac{\pi f C}{2L}}. \end{aligned} \tag{9}$$

Thus, the *gain in dB is inversely proportional to the square root of frequency* due to $R(f)$ which, in turn, is proportional to $\sqrt{f}$. This $\sqrt{f}$ response should not be confused with a pink-noise $\sqrt{f}$ response (in the pink-noise case, spectral density falls off at -10 dB/decade). Specifically, consider the curves shown in Fig. 3 for an ideal $\sqrt{f}$ cable and a first-order lowpass filter where the 3 dB frequency of each is at a normalized frequency of 1 Hz. In the case of a cable, the gain falls off slowly when the attenuation is small and then rapidly as the attenuation increases.

In terms of the cable's transfer-function phase, combining (4) and (8) we see that most of the phase is linear (which is a result of delay through the length of the cable) while a nonlinear component is proportional to the square-root of frequency.

Finally, using the typical primary constants above, we can find the gain for a typical twisted-pair cable to be approximately

$$H_{\mathrm{dB}}(d, f) \approx -0.02\, d\sqrt{f} \tag{10}$$

where d is in km and f is in hertz. Note that (10) is only useful in estimating the loss in a typical cable as it depends on the primary constants of the cable which in turn depend

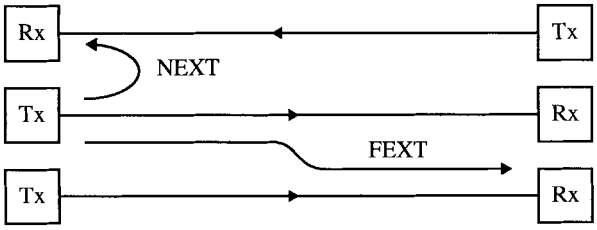

Fig. 4. Illustration of NEXT and FEXT.

on the physical construction of the cable (such as wire gauge, twist lengths, insulation type and thickness, etc.).

C. Crosstalk

As shown in Fig. 4, there are two main classes of crosstalk—near-end crosstalk (NEXT) and far-end crosstalk (FEXT). Far-end crosstalk is due to the leakage of transmitted signals onto a nearby wire pair and interfering with a locally transmitted signal. Thus, in the case of FEXT, the interfering crosstalk and the desired transmitted signals are both attenuated by the cable length. Near-end crosstalk is a result of large transmitted signals immediately leaking onto a nearby wire pair where small received signals from a distant transmitter are present (the received signals are attenuated due to the length of the cable). Thus, when present, NEXT almost always dominates FEXT in terms of limiting performance. In cases where the crosstalk interferer source is available, performance can be restored by cancelling NEXT similar to echo cancellation.

A reasonable model for the squared-gain of NEXT is given by

$$|H_N(j\omega)|^2 = K_N|\omega|^{1.5} \tag{11}$$

where K_N is a constant determined by physical parameters of the cables [4]. It should be mentioned here that a measured crosstalk gain will actually have many peaks and nulls over the frequency spectrum. However, the above relation is reasonable for modeling the envelope of the peaks in crosstalk gain.

Note that the crosstalk coupling increases with frequency. Since the cable transfer-function decreases with frequency, there will be some frequency value where the interference power equals the signal power. Thus, transmitting signal power above this frequency value is of little use.

III. Applications

A. HDSL

HDSL is a transmission standard for full-duplex delivery of 1.544 Mb/s over long lengths (up to 4.5 km) of existing telephone-loop UTP cables. To accomplish this goal, two pairs of cables are typically used along with a four-level pulse-amplitude modulation (PAM) code referred to as "2B1Q" (two bits—one quaternary). In other words, one of four levels are sent corresponding to 2 b of information. As a result, each cable must operate at a rate of 386 kS/s (kilo-Symbols per second) in full-duplex mode (actual transmission rate is 392 kS/s with overhead).

Detailed cable modeling in the HDSL environment was presented in [4] where it was shown that the $\sqrt{f}$ attenuation model is valid at both low and high frequencies and has a transition region where the attenuation falls off more slowly. Making use of (10) and letting $d = 4.5$ km, we see that a 200 kHz signal is attenuated by roughly 40 dB. Thus, the high-frequency portion of a $5\,V_{\rm pp}$ transmit signal would appear on the receive end as a $50\,{\rm m}V_{\rm pp}$ signal! As a result, effective echo cancellation becomes one of the main design challenges in HDSL.

When receive signals are small in relation to the transmit signals, it is not uncommon for the linearity of the echo path to limit performance of the overall system. Such a limitation occurs because the echo cancellation circuitry cannot typically deal with nonlinearities and/or noise in the echo path. For example, when the receive signal is attenuated by 40 dB, if the 2–4 wire hybrid performs only 6 dB echo attenuation (a typical number due to the large load variations in HDSL applications), the unwanted transmit signal is 34 dB above the desired receive signal at the input of the A/D converter. Since the echo canceler can eliminate all of the linear portion of the transmit signal (but not nonlinearities or noise), then to ensure the residual transmit signal is 40 dB below the desired receive signal after echo cancellation, the linearity and noise in the transmit path should be better than 74 dB. Referring to Fig. 1, the blocks that need to meet this 74 dB requirement are the transmit D/A, the line-driver, the 2–4 wire hybrid, and the receive A/D. This high linearity and noise requirement in the echo path is one of the most demanding aspects in realizing a fully-integrated HDSL transceiver.

Once the transmit echo signal is removed, the performance limit is mainly due to NEXT since HDSL signals are often sent in cable bundles where many nearby large transmit signals can couple into the received signal path.

B. Fast-Ethernet

Presently, ethernet accounts for approximately 85% of all local-area-network (LAN) implementations. For modest data rates, 10 Mb/s ethernet is realized over UTP cabling using a standard known as 10 Base-T. Within this LAN environment, there are two main types of UTP cables available—CAT3 and CAT5. The 10 Base-T standard operates well over both types of cabling.

CAT3 cables have poor attenuation and crosstalk characteristics. The attenuation response of 100 m CAT3 cables at 20°C is reasonably well modeled by

$$H_{\rm dB}(f) = 2.32\sqrt{f} + 0.238f \tag{12}$$

where f is in MHz [5]. Note that there is a term which is proportional to f in this model, and it should be taken into account during equalization. This f term is a result of a finite conductance value, G, due to the use of poor insulators. CAT5 cables have better attenuation and crosstalk characteristics. A reasonable model for of 100 m of CAT5 cable at 20°C is

$$H_{\rm dB}(f) = 1.967\sqrt{f} + 0.023f + \frac{0.05}{\sqrt{f}} \tag{13}$$

with f in MHz. Note that this type of cabling follows the $\sqrt{f}$ attenuation curve much more closely.

Fast-ethernet refers to a family of 100 Base-T standards intended to upgrade ethernet networks to 100 Mb/s over 100 m of CAT3 and/or CAT5 cabling. Due to the two main types of UTP cabling available, three separate substandards have emerged within the 100 Base-T standard.

100 Base-T4: 100 Base-T makes use of four pairs of CAT3 (or better) cables where two pairs are used half-duplex while the remaining two pairs are dedicated into one transmit and one receive pair. For example, if pairs 1–4 are available, pairs 1, 2, and 3 are used for transmission in one direction while pairs 1, 2, and 4 are used for transmission in the other direction. To reduce radiated emissions, eight binary bits are coded into six ternary (i.e., one of three levels) PAM outputs. As a result, the data rate of 100 Mb/s is achieved by sending 33 Mb/s over three pairs with a signaling rate of only 25 MS/s on each pair due to the 6/8 coding.

100 Base-TX: 100 Base-TX makes use of two pairs of CAT5 wiring each dedicated to either transmit or receive. As in 100 Base-T4, a three-level PAM code is used to reduce radiated emissions.

100 Base-T2: 100 Base-T2 makes use of two pairs of CAT3 (or better) cables each full-duplex through the use of 2–4 wire hybrids. However, to accommodate the difficult radiated emission standards, a 5×5 code is used where one of five PAM levels will be transmitted on each pair. The received constellation point is found by determining the received levels on both pairs. For 100 Mb/s, the transmission rate over each pair is 25 MS/s since one 5×5 symbol is sent for every 4 b (i.e., effectively 2 b on each pair) with nine of the 25 possible symbols used for signaling or unused.

To determine typical attenuation values for fast-ethernet applications, we make use of (10) and letting $d = 0.1$ km, we see that 12.5 and 20 MHz signals are attenuated by about 7 and 9 dB, respectively. These values might be typical attenuation numbers for CAT5 wire as seen also through the use of (13). Similar values in CAT3 wiring can be found from (12) and correspond to 11 and 15 dB, respectively.

IV. Main Building Blocks

A. Transmit D/A

The transmit D/A block typically includes filtering to shape the transmit spectrum. This filtering can be realized in both the digital and analog domains as shown in Fig. 5. Here, the digital transmit signal is upsampled and passed through a digital filter (these two operations are usually combined through the use of "poly-phase" filtering [6]). Such an approach requires a faster D/A converter but it reduces the requirements for the analog filter.

One of the challenges in HDSL realizations is to ensure better than 72 dB linearity in the overall echo path. Thus, it is important that the D/A converter have better than 12-b linearity, which is difficult to achieve with component matching. To overcome this challenge, oversampled 1-b D/A converters have been used as is commonly done in audio applications.

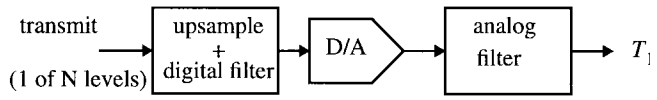

Fig. 5. Typical transmit D/A block.

For example, in [7] and [8], the transmit signal is upsampled by 16 times, applied to a 48-tap finite impulse response (FIR) digital filter and upsampled once more by four (for an overall oversampling rate of 64 times). This digital signal is then applied to a delta–sigma digital-to-audio converter (DAC) (with one-bit current output) to obtain an overall linearity and noise performance of at least 12 b. The use of an oversampled 1-b DAC eliminates the need for accurate on-chip component matching (to the 12-b level). A similar 1-b oversampled DAC approach was used in [9].

At higher speeds, such as in fast-ethernet applications, the linearity and noise requirements of the echo path are not nearly so severe. A typical requirement might be around the 40 dB level. For example, a proposed transmit path for 100 Base-T2 includes upsampling by three, using a 75 MHz 4-b DAC, and following this with a third-order lowpass filter. The design challenges here are relatively modest except for high-speed operation. One approach for adjusting a continuous-time pulse-shaping filter at high speeds makes use of comparator outputs at specific sampling instances and is presented in [10].

B. Line Driver

As seen in Fig. 1, a line-driver is used to supply the necessary drive current to the twisted-pair cable. These line-drivers are commonly realized as a voltage buffer having a low-impedance output.

As with other portions of the echo path, HDSL line-drivers need to have at least a 72 dB performance. Such a requirement for a power amplifier makes this block one of the more difficult challenges in the analog interface. Various circuit architectures for linear line-drivers were developed for integrated services digital network (ISDN) applications [11]–[13]. However, while the frequency requirement of ISDN drivers is only 40 kHz, it is nearly five times higher for HDSL with a similar linearity requirement.

To achieve large output swings and low distortion, modern line-drivers make use of fully differential structures. One possible architecture is shown in Fig. 6 [13]. Here, four large transistors are used and connected in a common-source configuration. As well, two feedback loops are used to obtain high linearity, speed, and output swings.

An alternative is to not place immediate feedback around the four internal amplifiers and instead rely on the main feedback path from input to output. For example, an HDSL line driver is shown in Fig. 7 where off-chip resistive feedback was used to maintain high linearity. In this case, a more elaborate compensation scheme was necessary such as the nested-Miller approach [14]. This circuit was realized in a 0.6-μm CMOS process and had a unity-gain frequency near 200 MHz. With a 10 kHz input signal, the total harmonic distortion (THD) of this line driver was measured to be better than 75 dB.

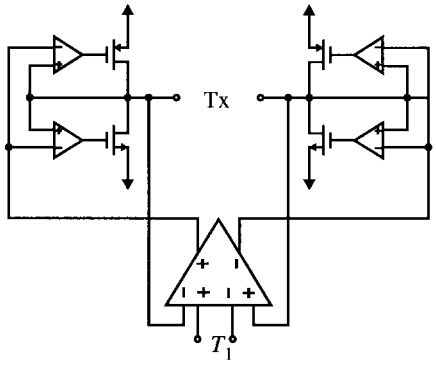

Fig. 6. A possible fully differential line driver.

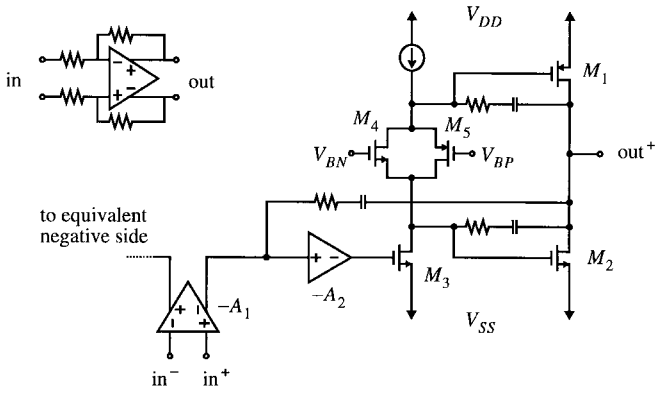

Fig. 7. An HDSL line driver with nested Miller compensation.

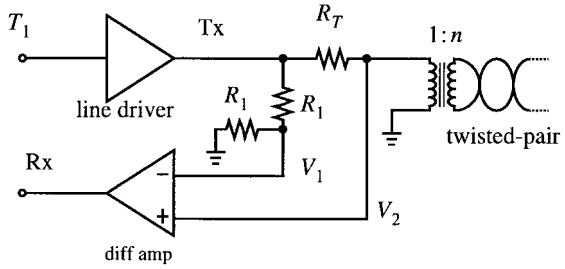

Fig. 8. A single-ended 2–4 wire hybrid interface circuit.

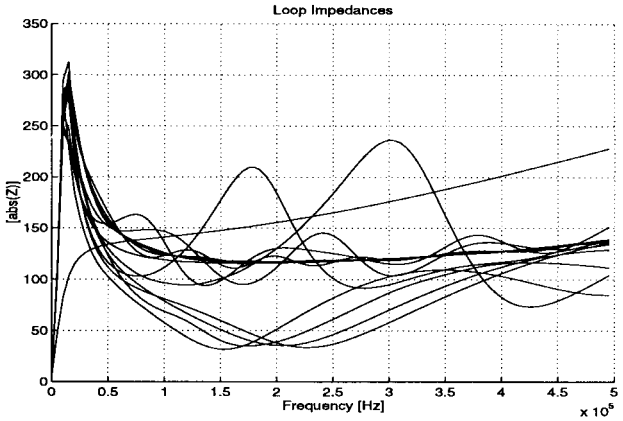

Fig. 9. Typical impedances seen on HDSL lines (transformer included).

A challenge in both these types of designs is ensuring that quiescent current is well controlled as well as the usual need for a common-mode feedback circuit. Note that it is imperative that the distortion of the final output transistor stage be modest since the open-loop gain will be reduced at high frequencies and will therefore not help much in improving linearity. For this reason, the last stage is likely biased as class-AB with a relatively large quiescent current. The class-AB biasing shown in Fig. 7 is accomplished through the use of bias voltages V_{BN}, V_{BP} and transistors M_4, M_5 as discussed in [15].

In fast-ethernet applications, the linearity of the line-drivers is not as stringent as in the HDSL application since received signal levels are only around 10 dB below transmit levels. In addition, the echo cancellation of the 2–4 wire hybrids is improved since the impedance seen by the cable is more controlled compared to the HDSL case. As a result, line driver linearity need only be around the 35 dB level. However, due to the high speeds involved and relatively low impedances that must be driven, many fast-ethernet implementations make use of bipolar transistors through the use of a BiCMOS technology.

C. 2–4 Wire Hybrids

The use of 2–4 wire hybrids allows one to transmit and receive simultaneously over a single cable (i.e., full-duplex transmission). However, hybrids are also commonly used even when two separate cables are available to reduce the frequency of the transmitted signal on each pair. Such a reduction in signal frequency is often important to meet emission requirements and can be particularly beneficial for longer cable lengths. Specifically, when longer cables are used, the received signals radiate little energy, and therefore, one is most concerned with reducing the high-frequency content of the transmitted signal. The main disadvantage of a full-duplex scheme is the need for excellent echo cancellation.

A possible single-ended hybrid configuration is shown in Fig. 8. Here, assuming the line impedance seen through the transformer equals R_T, then $V_2 = V_1$, and the transmit signal is cancelled. However, due to line impedance variations, the echo cancellation is only moderate. For example, a plot of various line impedances (including the transformer) for HDSL applications are shown in Fig. 9. Note that between 5 and 100 kHz, the impedance falls off at roughly $\sqrt{f}$ as expected from (3). Above 100 kHz, the majority of impedances are around 100 Ω, but there are lines which go as low as 40 Ω as well as some that go over 200 Ω. These large line impedance variations are due to unterminated bridge-taps and one would not expect as large a variation in fast-ethernet applications. Note that due to the transformer, there is a zero at dc and pole around 5 kHz. This low frequency pole results in a long impulse response in the echo path. To shorten the echo impulse response, the pole can be placed at higher frequencies (either through the use of a different transformer or digitally after the A/D converter), but then channel bandwidth is reduced.

Other variations on this hybrid circuit are to eliminate the R_1 network altogether and rely on the digital echo canceler to remove the transmit signal. Such an approach would require more dynamic range in the receive path (i.e., more bits on the A/D converter). Alternatively, the R_1 divide-by-two circuit might be made more complex so that better echo cancellation

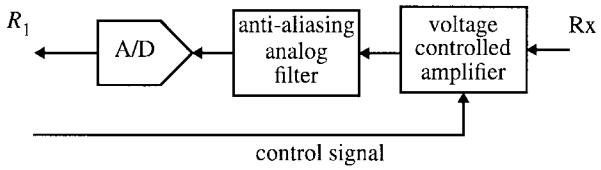

Fig. 10. Typical receive A/D block.

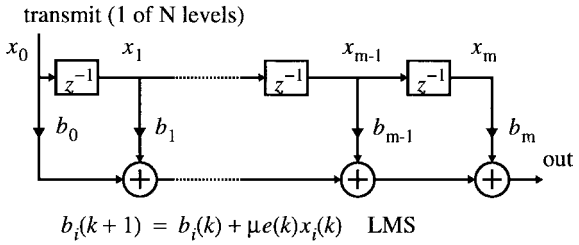

Fig. 11. An adaptive transversal FIR filter and LMS algorithm.

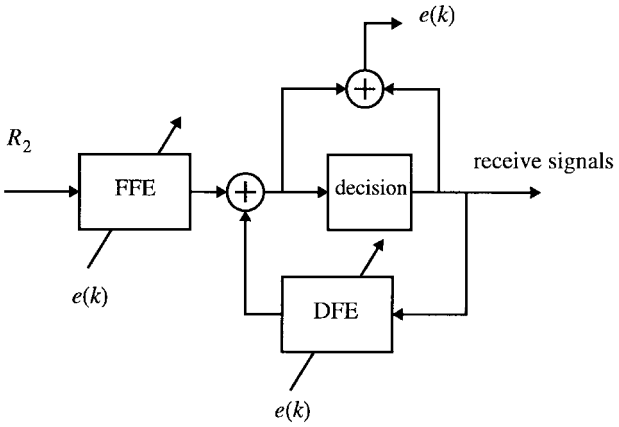

Fig. 12. A typical equalizer with feed-forward and feedback equalization.

occurs, thereby reducing the requirements on the receive path. Specifically, the two resistor R_1 circuit can be replaced with a circuit that better matches a nominal transfer-function from Tx to V_2. Finally, fully differential realizations of this hybrid circuit are also commonly used.

D. Receive A/D

A typical receive A/D block is shown in Fig. 10. Note that to reduce analog complexity, the control signal for the voltage controlled amplifier is often derived from a digital signal after A/D conversion [as opposed to an all analog automatic gain control (AGC)].

The requirements of the anti-aliasing filter depend on the sample-rate of the A/D converter. Thus, in lower speed applications, such as HDSL, oversampling can be used to ease these requirements at the expense of more digital filtering after A/D conversion. In fast-ethernet applications, this analog filter is most likely continuous-time due to its high speed and moderate linearity requirements. For example, a proposed receive block for the 100 Base-T2 standard includes a fifth-order continuous-time lowpass filter here with a 6-b A/D sampled at three times the symbol-rate (i.e., 75 MHz). One of the challenges here is to keep the size and power consumption of this fast A/D converter low.

E. Echo Cancellation

Echo cancellation is typically realized as an adaptive transversal finite impulse response (FIR) filter making use of the least mean square (LMS) algorithm as shown in Fig. 11. Note that the input to the transversal filter is the transmit symbols, and therefore, full multiplications are not needed within the filter or in the LMS algorithm due to the limited number of transmit symbols. Similarly, the number of shift-registers bits needed in the delay chain is also reduced.

Due to the high dynamic range requirement and low frequency pole (because of transformer coupling) in an HDSL echo path, the impulse response of the echo path in HDSL is quite long and requires approximately a 120-tap FIR adaptive filter [16]. Besides requiring many multiplications and additions, another disadvantage of a large FIR filter is that the coefficient bit size increases as the number of taps increases. Specifically, it can be shown that because of coefficient quantization, every four times increase in tap length requires another bit of coefficient accuracy. Thus, although an HDSL application needs only 13-b dynamic range, coefficient accuracy in the echo canceler might be around 20 b. As a result, this echo cancellation circuitry can consume a large amount of silicon area.

In fast-ethernet applications, there are lower dynamic range requirements because of larger receive signals, and thus a typical echo canceler might be around 30 taps. Also, it is sometimes possible to add some NEXT cancellation when one is transmitting on two (or more) cables in the same bundle. Specifically, if one is transmitting on cables 1 and 2, then the NEXT from cable 1 onto cable 2 can be cancelled by applying the transmit signal from cable 1 into the receiver circuitry of cable 2. Similarly, the NEXT from cable 2 to cable 1 can also be cancelled. Of course, the NEXT due to unrelated cables in the same bundle remains.

F. Equalization

Assuming digital equalization is performed, it is often accomplished through the use of a feed-forward equalizer (FFE) and decision-feedback equalizer (DFE) as shown in Fig. 12. While an FFE alone could be used, it is usually arranged that the FFE removes precursor ISI, while the DFE removes postcursor ISI.

One advantage of using a DFE is reduced noise enhancement since the input to the DFE is noise free (it is assumed to be the correct received signals). However, error propagation reduces the effectiveness of the DFE since if a wrong decision is made, this wrong decision remains within the tapped-delay line of the DFE until it is flushed out. Another advantage of using a DFE is that the multiplications within the DFE are trivial since the received signals take on a limited number of values (say one of four levels). In the case of the FFE, full multiplications are required. As a result, one would like to use a smaller number of taps in the FFE in comparison to the DFE. Fortunately, such an arrangement is often possible since the channel tends to have a longer tail following the main peak

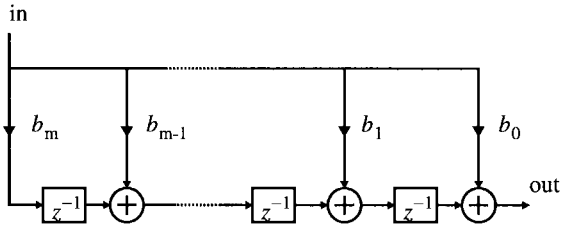

Fig. 13. Transposed transversal FIR filter.

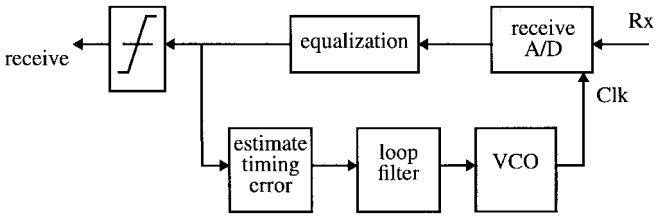

Fig. 14. A clock recovery system (echo cancellation not shown).

than preceding it. For example, one HDSL approach made use of a 19-tap FFE and a 128-tap DFE [9].

One issue that can make the DFE difficult to implement is that there should be low latency through the DFE, otherwise the first few postcursors cannot be removed. Such a low latency requirement can be difficult in a standard transversal filter where the output is formed as the sum of many signals. For example, in an m tap transversal filter as shown in Fig. 11, m signals need to be combined within one clock cycle. To overcome this large addition requirement, one can use a transposed transversal filter as shown in Fig. 13 [17]. Here, the input signal is applied to many multipliers in parallel, but there is a delay block between each adder. Thus, each adder has only two inputs and there is a delay between adders. It should be noted that the gradient signals are the same for the filters of Figs. 11 and 13, and therefore, a separate delay line is needed in the case of the transposed filter to realize the gradient signals. Another disadvantage of the transposed structure shown in Fig. 13 is that the delay blocks need to be at least as large as the coefficient widths (say 20-b wide) as opposed to only 2-b wide in the transversal filter of Fig. 11.

One common alternative to the above equalizer architecture is to use a fractionally-spaced FFE such that both matched filtering and equalization are performed within the FFE [1]. A fractionally-spaced FFE operates at a higher rate than the data rate (say three times) and can be efficiently implemented as a polyphase filter [6], assuming downsampling is performed immediately following the FFE. Another variation is to include an error-prediction filter to aid in adaptation [7].

While digital equalization is common at HDSL rates, analog equalization is a serious alternative for applications such as fast-ethernet. The advantages of analog equalization are reduced size and power. As well, in applications such as 100 Base-Tx, the cable transfer-function is reasonably well modeled with little variation once the cable length is known. As a result, a common analog equalization approach is to estimate the cable length based on the size of the received signal and thereby tune the analog filter appropriately. More complex analog equalizers have not found common usage mainly due to the difficulty in designing complex analog systems as well as their susceptibility to dc offsets [18].

G. Clock Recovery

A typical clock recovery arrangement is shown in Fig. 14. If the complexities of the equalizer and echo canceler are relatively modest, the A/D can be operated at twice the symbol rate such that an early–late gate approach is used [1], [19]. Specifically, the output of the equalizer is observed before and after the decision instant by $T/2$ (where T is the symbol rate) and the VCO is adjusted to make these two intermediate values equal. However, this approach requires that the echo canceler and the equalizer run at twice the symbol rate.

In order that the rate of the equalizer or echo canceler not be increased, baud-rate timing recovery is often used [1], [20]. This approach makes use of decision-directed timing adjustments that are performed in the digital domain and then indicate to the VCO whether the phase should be increased or decreased.

The VCO is often a frequency synthesizer that should display low jitter performance. While past approaches have relied on voltage-controlled crystal oscillators to maintain low jitter, a technique which is quickly gaining in popularity is to make a digitally controlled frequency synthesizer using oversampling techniques as shown in Fig. 15 [21], [22]. Here, a divide-by-N counter is set to one of three integer values by applying F, which can be precise to many bits, to an oversampled delta–sigma modulator and quantizing its output. As a result, the frequency accuracy of the VCO can be set quite precisely and the jitter noise is mostly at high-frequencies, allowing the loop filter to attenuate it significantly. An HDSL transceiver using this scheme illustrates the performance which can be achieved. This circuit, with $f_{xt} = 6.272$ MHz, $M = 1$, $k = 32$, and $P = 8$ has an rms jitter under 100 ps.

Finally, it should be mentioned that low-frequency clock jitter is often attenuated by the addition of a circuit known as an "elastic buffer." Specifically, when the input clock rate varies slowly around an average clock rate, it is often desirable that the output clock rate of the receiver is maintained at the average rate (in effect, attenuating the low-frequency clock jitter to any remaining circuits). However, while the average of the input and output clock rates must be equal, instantaneous clock differences are tolerated through the use of a FIFO buffer which remains approximately half full. For example, when the input rate is greater than the output rate, the receiver accumulates extra data bits in the buffer. On the other hand, when the input rate is less than the output rate, the receiver must have extra bits in the buffer to maintain a constant output rate. On average, the buffer is half full through the use of a feedback circuit which slowly adjusts the output clock rate (slower than the low-frequency jitter).

V. Future Challenges

While HDSL is presently delivered over two twisted-pair cables, there is much interest in single pair implementations. Realizing single-pair HDSL will effectively double the symbol

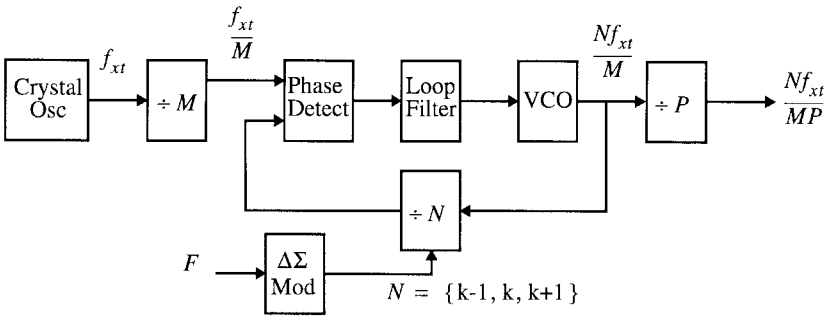

Fig. 15. A fractional-N frequency synthesizer.

rate unless a new signaling scheme is adopted. In the fast-ethernet area, vendors are currently looking at designing 100 Base-T2 chipsets that are cost competitive with TX and T4 solutions. As well, there is always interest in looking at other modulation schemes [23]–[25]. Finally, there is the usual push toward smaller geometry processes that allow more digital circuitry but will most likely challenge analog designers (such as only a 3.3 V power-supply voltage being available).

A. Line-Drivers

In HDSL applications, maintaining the linearity of line-drivers will continue to challenge circuit designers. Furthermore, as power-supply voltage levels decrease with smaller device dimensions, producing the necessary output line voltage will be difficult. One method to allow lower voltage levels from the line-driver is to adjust the turns ratio of the line-transformer. Unfortunately, this also lowers the impedance that must be driven by the line-driver. Another approach is to use a current line-driver as shown in Fig. 16. Here, the line-driver approximates an ideal current source and needs only supply half the voltage as the approach shown in Fig. 8. However, this current line-driver must source twice as much current, and it may be difficult to maintain the same linearity as the voltage approach.

B. 2–4 Wire Hybrids

If 2–4 wire hybrids can be realized with less echo return-loss, the linearity requirement of the echo path can be reduced (i.e., the D/A, line-driver, and A/D converter). Such extra echo cancellation may be necessary to realize single-pair HDSL. However, as the trend is toward more integration with less analog and discrete components, it is likely that this extra cancellation might cost extra A/D or D/A converters. One such hybrid echo cancellation circuitry is discussed in [26].

C. Echo Cancellation

As discussed above, the echo cancellation circuitry in HDSL applications can require over a 120-tap FIR filter. For higher data rates, this already long filter length will further increase. As noted in [16], there is likely some benefit of creating hybrid FIR/IIR-type echo cancellation circuits to reduce overall complexity. Thus, a challenge for circuit designers will be to realize effective adaptive IIR filtering.

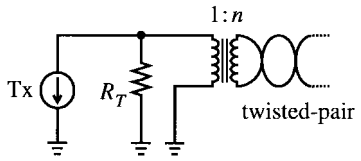

Fig. 16. A current-driven line-driver.

Another interesting avenue of research is to realize efficient nonlinear echo cancellation circuits such that the linearity requirements of the echo path can be relaxed. This challenge is complicated by the fact that the nonlinearities being introduced into the echo path are not typically memoryless and are therefore better modeled using Volterra series expansions rather than Taylor series.

D. Equalization

It is likely that digital equalization will be used on channels having modest data rates. To improve performance, it is sometimes advantageous to equalize to a lowpass response, thereby treating the channel as a type of partial-response system (this approach is now adopted by many disk-drive manufactures and is known as partial-response maximum likelihood (PRML) [27]). By reducing the channel boost at high frequencies, less noise is amplified. However, now some type of maximum likelihood detector is needed and is most likely suboptimal for complexity reasons [28]. While this maximum likelihood detection approach can ideally achieve an extra 3 dB performance, other imperfections in the receive signal and modeling typically limit the improvement to smaller values.

As mentioned above, for high-speed channels such as fast-ethernet, analog equalization is a serious alternative. The two main advantages of analog equalization are lower power and less silicon area. The main challenge in analog equalization remains dealing with dc offsets [18] as well as realizing efficient programmable filter structures.

E. Clock Recovery

Jitter is always an important performance criteria of clock recovery. While most present clock recovery techniques adjust the sampling instance, it is also possible to perform fully digital clock recovery by resampling the digital signal (in other words, use sample-rate conversion techniques [6]). While the complexity of digitally resampling is usually considered too

costly, there are advantages to be gained such as running the A/D converter off a crystal controlled clock having extremely low jitter. Another clock recovery challenge is that of tracking high-frequency jitter since the latency of the clock recovery loop often includes equalization and therefore has limited tracking speed. This latency can be shortened by not placing the equalizer and any extra filtering in the clock recovery loop but implies that either an oversampled or analog equalizer is used. The clock-recovery can then be accomplished by digital resampling at the cost of extra circuit complexity.

VI. Conclusion

This paper discussed modern architectures for realizing high-speed data transmission over twisted-pair cables and addressed two applications in particular—HDSL and fast-ethernet. With HDSL, the echo path linearity is seen to be one of the most challenging system requirements. In the fast-ethernet application, keeping the cost and power consumption low is perhaps the biggest challenge. However, with the wide variety of applications requiring data transmission over twisted-pair channels, perhaps the biggest challenge is in deciding which applications to pursue.

References

[1] E. A. Lee and D. G. Messerschmitt, *Digital Communication.* Norwell, MA: Kluwer, 1994.

[2] W. D. Reeve, *Subscriber Loop Signalling and Transmission Handbook: Digital.* Piscataway, NJ: IEEE Press, 1995.

[3] E. C. Jordan and K. G. Balmain, *Electromagnetic Waves and Radiating Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1968, p. 563.

[4] J. J. Werner, "The HDSL environment," *IEEE J. Select. Areas Commun.,* vol. 9, pp. 785–800, Aug. 1991.

[5] G. H. Im *et al.,* "51.84 Mb/s 16-CAP ATM LAN standard," *IEEE J. Select. Areas Commun.,* vol. 13, pp. 620–632, May 1995.

[6] P. P. Vaidyanathan, *Multirate Systems and Filter Banks.* Englewood Cliffs, NJ: Prentice-Hall, 1993.

[7] D. Essig *et al.,* "An HDSL chipset for high-speed transmission over copper," in *Proc. IEEE Int. Conf. Communications,* New Orleans, 1994, pp. 461–465.

[8] *HDSL Data Book.* Brooktree Corp., San Diego, CA, 1995.

[9] M. A. Kuczynski *et al.,* "A 1 Mb/s digital subscriber line transceiver signal processor," in *IEEE Int. Solid-State Circuits Conf.,* San Francisco, CA, Feb. 1993, pp. 26–27.

[10] A. Shoval, W. M. Snelgrove, and D. A. Johns, "A 100 Mb/s BiCMOS adaptive pulse-shaping filter," *IEEE J. Select. Areas Commun.,* special issue on *Copper Wire Access Technologies for High Performance Networks,* in press.

[11] F. N. L. Op'T Eynde, P. F. M. Ampe, L. Verdeyen, and W. M. C. Sansen, "A CMOS large-swing low-distortion three-stage class AB power amplifier," *IEEE J. Solid-State Circuits,* vol. 25, pp. 265–273, Feb. 1990.

[12] L. Tomasini, A. Gola, and R. Castello, "A fully differential CMOS line driver for ISDN," *IEEE J. Solid-State Circuits,* vol. 25, pp. 546–554, Apr. 1990.

[13] H. Khorramabadi, "A CMOS line driver with 80-dB linearity for ISDN applications," *IEEE J. Solid-State Circuits,* vol. 27, pp. 539–544, Apr. 1992.

[14] R. G. H. Eschauzier, R. Hogervorst, and J. H. Huijsing, "A programmable 1.5 V CMOS class-AB operational amplifier with hybrid nested Miller compensation for 120 dB gain and 6 MHz UGF," *IEEE J. Solid-State Circuits,* vol. 29, pp. 1497–1504, Dec. 1994.

[15] R. Hogervorst, J. P. Tero, R. G. H. Eschauzier, and J. H. Huijsing, "A compact power-efficient 3 V CMOS rail-to-rail input/output operational amplifier for VLSI libraries," *IEEE J. Solid-State Circuits,* vol. 29, pp. 1505–1512, Dec. 1994.

[16] W. Y. Chen, J. L. Dixon, and D. L. Waring, "High bit rate digital subscriber line echo cancellation," *IEEE J. Select. Areas Commun.,* pp. 848–860, Aug. 1991.

[17] H. Samueli, B. Daneshrad, R. B. Joshi, and B. C. Wong, "A 64-tap CMOS echo canceller/decision-feedback equalizer for 2B1Q transceivers," *IEEE J. Select. Areas Commun.,* pp. 839–847, Aug. 1991.

[18] A. Shoval, D. A. Johns, and W. M. Snelgrove, "DC offset performance of four LMS adaptive algorithms," *IEEE Trans. Circuits Syst.—II: Analog and Digital Signal Processing,* vol. 42, pp. 176–185, Mar. 1995.

[19] A. Buchwald and K. W. Martin, *Integrated Fiber-Optic Receivers.* Norwell, MA: Kluwer, 1995.

[20] K. H. Mueller and M. Muller, "Timing recovery in digital synchronous data receivers," *IEEE Trans. Commun.,* vol. COMM-24, pp. 516–531, May 1976.

[21] B. Miller and R. J. Conley, "A multiple modulator fractional divider," *IEEE Trans. Instrum. Meas.,* vol. 40, pp. 578–583, June, 1991.

[22] T. A. D. Riley, M. A. Copeland, and T. A. Kwasniewski, "Delta–sigma modulation in fractional-N frequency synthesis," *IEEE J. Solid-State Circuits,* vol. 28, pp. 553–559, May 1993.

[23] G. H. Im and J. J. Werner, "Bandwidth-efficient digital transmission up to 155 Mb/s over unshielded twisted pair wiring," in *Proc. IEEE Int. Conf. Communications,* Geneva, 1993, pp. 1797–1803.

[24] G. Cherubini, S. Ölçer, and G. Ungerboeck, "A quaternary partial-response class-IV system for 125 Mb/s data transmission over unshielded twisted-pair cables," in *Proc. IEEE Int. Conf. Commun.,* Geneva, 1993, pp. 1814–1819.

[25] W. E. Stephens and T. C. Banwell, "155.52 Mb/s data transmission on category 5 cable plant," *IEEE Commun. Mag.,* pp. 62–69, Apr. 1995.

[26] D. Mueller and A. Kaelin, "A hybrid HDSL echo canceler," in *Int. Symp. Circuits Syst.,* Seattle, WA, May 1995.

[27] M. H. Shakiba, D. A. Johns, and K. W. Martin, "A 200 MHz 3.3 V BiCMOS class-IV partial response analog Viterbi decoder," in *Custom Integrated Circuits Conf.,* Santa Clara, CA, May, 1995, pp. 567–570.

[28] M. V. Eyaboglu and S. Qureshi, "Reduced-state sequence estimation with set partioning and decision feedback," *IEEE Trans. Commun.,* vol. 36, Jan. 1988.

David A. Johns received the B.A.Sc., M.A.Sc., and Ph.D. degrees from the University of Toronto, Canada, in 1980, 1983, and 1989, respectively.

From 1980 to 1981, he worked at Mitel Corp., Ottawa, Canada, while from 1983 to 1985 he was an analog IC designer at Pacific Microcircuits Ltd., Vancouver, Canada. His doctoral work focused on analog and digital adaptive filters including the development of an orthonormal structure for analog filters. In 1988, he was hired at the University of Toronto where he is currently an Associate Professor. He has ongoing research programs in the areas of analog integrated circuits, oversampling, and digital communication circuits resulting in more than 40 publications. He has been involved in numerous industrial short courses and has just completed a one-year research leave with Brooktree Corp. in San Diego, CA, in the area of high-speed data communications.

Dr. Johns is a past and present Associate Editor for IEEE Transactions on Circuits and Systems.

Daniel Essig received the B.S. degree in electrical engineering from Iowa State University, Ames, in 1980.

He is currently with the Brooktree Division of Rockwell Semiconductor Systems, San Diego, CA, where he designs communication IC's. His most recent design was a single-chip HDSL transceiver. Earlier at Brooktree he designed an NTSC/PAL video decoder. Previously he has developed programmable digital signal processors for Texas Instruments and Analog Devices. He started his electronics career at an early age tongue testing transistor radio batteries.