

DAC Quantization-Noise Cancellation in an Echo-Canceling Transceiver *

Paul Hurst¹ and Andy Norrell²

¹ Department of Electrical and Computer Engineering
University of California, Davis, CA USA

² Consultant, Nevada City, CA USA

Copyright IEEE 2007. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

Abstract— Data converters, in particular the transmit digital-to-analog converter (DAC), should not limit performance in a digital-communication transceiver. Typically, the number of DAC bits is chosen to be large enough so that the effect of DAC quantization noise on the local receiver is negligibly small. As described in this paper, the DAC quantization noise can be cancelled in a echo-canceling full-duplex transceiver. The proposed quantization-noise cancellation allows use of a simpler, lower resolution DAC. The quantization-noise cancellation concept is described, and simulation results are presented that demonstrate its operation.

Keywords: Echo cancellation, digital-to-analog converter, noise cancellation, analog front-end, digital communication.

I. INTRODUCTION

Transceivers for digital communications typically consist of digital-signal-processing circuits and an analog front-end (AFE). While the digital circuits in the DSP have been benefiting from shrinking transistor dimensions in advanced CMOS technologies, the analog circuits have not [1,2]. There are several ways to capitalize on the shrinking size of digital circuits. The amount of analog circuitry can be minimized (e.g., by moving analog filtering to the digital domain); the analog circuit blocks can be simplified (e.g., fewer bits in the data converters or lower order filters), resulting in small analog circuits; or digital calibration can be used to overcome the limitations of analog circuits. In all cases, extra digital circuitry is needed to perform functions that have been eliminated from the analog domain, to relax the requirements for the analog circuits, or to implement the calibration function.

In a full-duplex system, echo cancellation can allow bi-directional transmission over one channel using a fully or partially overlapping frequency band for both transmit signals [3]. In this paper, a modified transceiver architecture is proposed to simplify a key analog circuit block in an echo-canceling transceiver. The proposed transceiver adds an adaptive digital filter to cancel quantization noise associated with the digital-to-analog converter (DAC) in the transmit signal path. Using this quantization-noise cancellation block, the number of bits required in the transmit DAC can be significantly reduced.

The proposed approach is intended for medium to high-speed transceivers that use Nyquist DACs. For lower-speed systems, DAC quantization noise can be made low using a noise-shaping delta-sigma DAC.

Background is presented in the next section, followed by a description of the proposed idea in a pulse-amplitude-modulation (PAM) based system. Simulation results are presented that successfully demonstrate the proposed quantization-noise-cancellation concept.

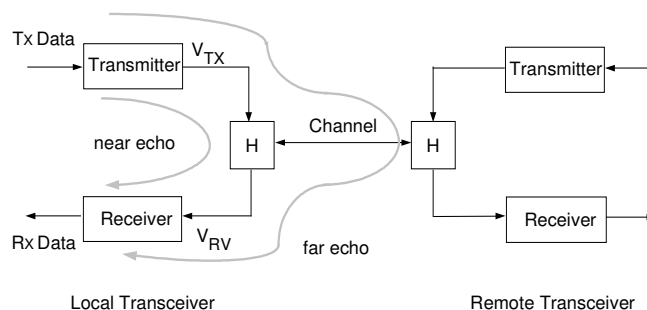


Fig. 1. Two transceivers in a full-duplex system. Shown are the transmitter, receiver and hybrid (H) blocks. The echo cancellers are not shown here. The near and far echo paths for the local transceiver are also shown.

II. BACKGROUND

Fig. 1 is a simplified block diagram of a full-duplex digital communication system. The left-hand (or local) transceiver transmits data to the right-hand transceiver over the two-wire channel while the right-hand (or remote) transceiver transmits over the same channel using an overlapping frequency band. Consider the local transceiver, which is shown in more detail in Fig. 2. The hybrid interfaces the transmit and receive circuitry to the two-wire channel [3,4]. Ideally, the hybrid passes the local transmit signal V_{TX} onto the channel, matches the channel impedance to avoid reflections, and isolates the receiver input signal V_{RV} from the local transmit signal V_{TX} . Ideally, the local receiver input V_{RV} consists only of signal from the remote transmitter.

* This research was funded in part by UC MICRO grant 06-204.

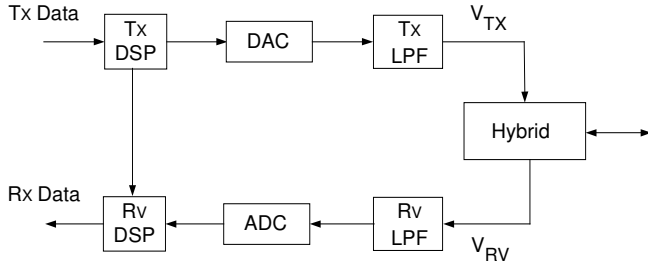


Fig. 2. Block diagram of a transceiver, showing some of the key blocks.

In practice, the hybrid is imperfect, and the receiver input V_{RV} consists of the signal sent by the remote transmitter, filtered by the channel, plus a filtered version of the signal from the local transmitter, called the echo. An echo canceller (EC) can be used to cancel this echo, allowing the back-end of the local receiver to extract the data from the remote transmitter [3,4].

The echo signal can be broken into two distinct components, which are shown in Fig. 1. The first component is near echo. Near echo stems from leakage of the local transmitter output into the local receiver input, due to imperfect operation of the local hybrid. The other echo component is far echo, which is a result of reflection of the local transmitted signal when it reaches an impedance discontinuity along or at the end of the channel. One measure of hybrid performance is trans-hybrid loss (THL). With the remote transmitter off in Fig. 1 [3,4],

$$\text{THL} = -20\log[V_{RV}(\text{rms})/V_{TX}(\text{rms})] \quad (1)$$

The EC is designed to remove both the near and far echo. Since the echo response is unknown, the EC is an adaptive finite-impulse-response (FIR) filter; and its coefficients are usually adapted using a least-mean square (LMS) algorithm [3].

A transceiver that uses an EC is shown in Fig. 3. The PAM Tx Data signal is upsampled and filtered by a digital filter (TxFilt); this filter can relax the requirements of the analog transmit LPF (TxLPF) and provide compromise pre-equalization for the average channel. The Q block is a digital quantizer, which takes an input sample of B_1 bits and produces an output sample with B_2 bits ($B_2 < B_1$) by truncating or rounding. A DAC produces an analog signal from the quantizer output, and its output is filtered by analog filter TxLPF. On the receive side, the input is low-pass filtered (RvLPF), digitized by the analog-to-digital converter (ADC), and then processed by the receive digital-signal processor (DSP) to extract the received data. (A variable-gain amplifier often precedes the ADC, and a line driver typically follows the TxLPF and drives the channel. These blocks are not shown in Fig. 3 for simplicity.)

In the transmit path, assume the TxFilt outputs a B_1 -bit signal that must be quantized to B_2 bits for a B_2 -bit DAC ($B_2 < B_1$). Due to this quantization, quantization error or 'quantization noise' of $B_1 - B_2$ bits is present in the quantizer

output, in addition to the desired transmit signal. This quantization noise is undesirable and could limit transceiver performance.

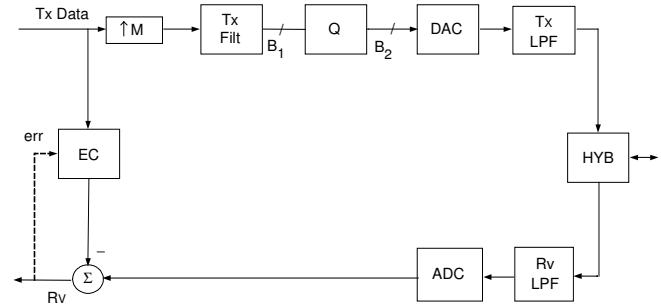


Fig. 3. Block diagram of a transceiver using the conventional EC approach.

Assuming perfect echo cancellation and ignoring ADC quantization, an important implementation-related noise source in the local receiver is local DAC quantization noise. In one approach (the conventional EC approach in Fig. 3), the number of bits in the local transmit DAC is chosen large enough so that the echo of DAC quantization noise at Rv in the local receiver is well below noise from other sources, such as noise from the channel. For example, assume the desired SNR at Rv in Fig. 3 is 25 dB and any local DAC quantization noise there should be 15 dB below the noise from other sources. Then a signal-to-quantization-noise ratio (SQNR) of 40 dB is desired at Rv. For simplicity, assume all analog blocks except the DAC in Fig. 3 are ideal with unity gain, and assume a perfect EC. If the peak-to-rms ratio of the remote transmit signal is 12 dB and the remote DAC full-scale is 0 dBV, the remote transmitter can output a signal at -12 dBV. If the channel attenuation is 36 dB, the signal at Rv from the remote transmitter will be -48 dBV. To achieve a SQNR of 40 dB, the local DAC quantization noise at Rv must have an rms value of -88 dBV. If the trans-hybrid loss is 12 dB and the peak-to-rms ratio is 12 dB for the local transmit signal, then a DAC with $B_2 = 12$ bits is needed. ($-88\text{dBV} = -\text{THL} + [\text{rms of quantization noise in dBV}] = -12\text{dB} + [(\text{rms of DAC output signal in dBV}) - \{\text{DAC SQNR in dB}\}] = -12\text{dB} + [-12\text{dBV} - \{6.02(B_2 - 1) - 1.24\} \text{dB}]$, so $B_2 \approx 12$. The SQNR formula is from [5].) Such a high resolution DAC at a high conversion rate in a low-voltage CMOS process would be a design challenge, and it might consume significant IC area and power.

An advantage of this approach is that the EC input is baud-rate samples of the PAM transmit signal. Each sample often consists of only a few bits, (e.g., 4 bits/symbol for HDLS2 [6]; 3 bits/symbol for 1Gb/s Ethernet [7]), simplifying the implementation of the delay line and multipliers in the EC. The key disadvantage is that many bits (e.g. 10-14 bits [8,9,10]) are needed in the DAC.

An alternative approach is to build an EC that operates on the quantized DAC signal. This approach is shown in Fig. 4. Here the input of the EC is taken after the quantizer

Q; so the EC input includes the quantization noise. Therefore the EC can cancel the echo of the transmitted signal, which includes the DAC quantization noise. Hence, the number of bits in the transmit DAC (B_3) can be determined solely by the requirements of the remote receiver. Considering only DAC quantization noise, the SQNR at Rv in the remote receiver will equal the SQNR at the local DAC output. To achieve SQNR=40 dB at the DAC output with a 12dB peak-to-rms ratio of the transmit signal, an 8-bit DAC can be used (using SQNR = $\{6.02(B_3 - 1) - 1.24\}$ dB [5].) This is 4 bits less than needed in the previous example for the conventional EC approach in Fig. 3.




Fig. 4. Simplified transceiver with the EC operating on the quantized DAC input signal.

An advantage of this approach is a simpler DAC, compared to the conventional approach. However, the key drawback is that the EC complexity may be high. The EC can require a long FIR filter, up to hundreds of taps [11,7], to cancel both the near and far echo signals. A long EC with 8-bit input samples requires significantly more circuitry (requiring more area and power) in the delay line and the multipliers than in the conventional case, where the EC input contains only a few bits. Also, the input to the EC in Fig. 4 is at a higher sampling rate than in Fig. 5, leading to more computation. Finally, the EC convergence here may be adversely affected by the colored spectrum of its input, which has been filtered by TxFilter.

III. PROPOSED TRANSCEIVER USING A QUANTIZATION NOISE CANCELLER

The DAC quantization noise $q[n]$ can be calculated in the digital domain by computing the difference between the input and output of the quantizer Q that precedes the DAC, as shown in Fig. 5. This quantization noise is present in the DAC output signal sent to the remote receiver and in the echo. The DAC quantization noise $q[n]$ can be cancelled using an adaptive noise-cancelling FIR filter, called the quantization-noise canceller (QNC), as shown in Fig. 5. For simplicity, assume the ADC is ideal and does not introduce quantization noise. If the transfer function of the QNC filter is equal to the transfer function from the DAC input to the ADC output in the local receiver, then the DAC quantization noise present in the ADC output will be completely cancelled by the QNC. In that case, the quantization noise in the local Rv signal will be due to the B_1 -bit quantization of the quantizer (Q) input. Assuming perfect QNC operation, the number of DAC bits (B_4) can be chosen solely based on the SQNR requirement of the remote

receiver. Using the values in the examples above, the required number of transmit DAC bits is 8.

The quantization noise from the local DAC is an undesired signal that appears in the local receive signal path due to nonzero echo. If the input to the quantizer is a 'busy' signal (as is desired), the quantization noise will be approximately white and uncorrelated with the quantizer input [5]. The QNC uses an adaptive FIR filter with input $q[n]$ to approximate the filtered version of $q[n]$ present at the ADC output. The QNC filter coefficients in Fig. 5 can be adapted using LMS adaptation [3]:

$$c_j[n+1] = c_j[n] + \mu_q \text{err}[n] q[n-j] \quad (2)$$

where $c_j[n]$ is the j^{th} QNC coefficient at time n , $\text{err}[n]$ is the error signal, which is also the input to the back-end of the receiver, and μ_q is the LMS update gain. The QNC and EC are adapted jointly, using the same error signal, $\text{err}[n]$ [3].




Fig. 5. Proposed transceiver that uses a QNC and an EC.

Two adaptive filters are needed here for the EC and QNC. At first, the complexity of having both an EC and QNC may seem high. However, in some applications the far echo of the DAC quantization may be small enough to ignore; then the QNC can be a simple, short-length filter. This situation is considered next.

In Fig. 5, the number of bits in the local transmit DAC is chosen so that the DAC quantization noise output is well below the acceptable noise floor at the input to the remote receiver. Assume the local and remote transmitters operate at the same output power and the far echo stems from impedance mismatch at the far end of the channel. Consider both the far echo at the local receiver input and the signal from the remote transmitter at the local receiver input (see Fig. 1). The signal from the remote transmitter traverses the channel once. The strongest far echo corresponds to one reflection of the local transmitted output after it traverses the channel, reflects at the far end, and then traverses the channel again before entering the local receiver.

If the channel attenuation is A dB and the return loss at the interface to the far-end transceiver is R dB, then at the channel input of the local hybrid, the transmit signal from the remote transmitter and the quantization noise from the remote transmit DAC have been attenuated by A dB by the channel. The largest far echo signal from the local transmitter is attenuated by $2A+R$ dB, due to traversing the channel twice and reflecting from the far end. As a result, the local DAC quantization noise in the far echo has been attenuated by $2A+R$ dB when it appears at the channel input

of the local hybrid. Therefore, the quantization noise in the far echo from the local DAC is $A+R$ dB smaller than the quantization noise from the remote DAC at the local hybrid input. With $A \geq 0$ dB and $R \geq 12$ dB (a reasonable value [4,12]), the local DAC quantization noise in the far echo is well below the DAC quantization noise from the remote transmitter; therefore, the far echo of the DAC quantization noise has little impact on the local receive signal. Hence, the QNC may only have to span the near echo, which can be done with a short FIR filter in the QNC in most systems.

The approach in Fig. 5 has several advantages: a simpler DAC (with fewer bits) than the conventional approach, a simple EC (with few input bits), and a QNC that will usually have few input bits and may only have to span the local echo. Also, if the TxFilter output is 'busy' [5], the input signals for the EC and QNC are uncorrelated and white, which is beneficial for convergence. However, the DAC linearity must be good enough to keep distortion of the transmit signal sufficiently low, as the QNC does not cancel DAC distortion in the receive signal. Also, the DAC circuit noise is not affected by the QNC.

IV. SIMULATION RESULTS

The system presented in Figs. 1 and 5 was simulated under various conditions. In all cases, referring to Fig. 2, random 8-level PAM (PAM8) transmit data (TxData) is up-sampled by $M=2$ and passed through TxFilter, which is a square-root raised-cosine filter [3] with 20% excess bandwidth cascaded with a first-order pre-emphasis filter that compensates for the roll-off of the average channel. The output of TxFilter and the input to quantizer Q is a 12-bit word ($B_1=12$). The channel and echo responses were determined from simulation and measurement of an Ethernet channel (CAT6 cable) and a passive hybrid. The channel consisted of two 5m patch cables in series with a 25m cable. The impulse response of the total (near + far) echo is plotted in Fig. 6. The near echo is about 16 samples long, and the far echo can be seen between samples 70 and 100.




Fig. 6. Impulse response of total (near and far) echo.

An ideal ADC is used in the receiver (no ADC quantization). The only noise source in the simulations is DAC quantization noise. The ADC and DAC operate at a sampling rate of 1600 MHz, which is twice the baud rate of

800 MHz (this is the proposed symbol rate for 10Gb/s Ethernet over copper [13]). The EC uses a poly-phase structure and spans 128 samples, and unless stated otherwise, the QNC FIR filter length is 128 samples. With these filter lengths, the EC and QNC span the total (near and far) echo. The EC and QNC use LMS adaptation [see (2)].

Fig. 7 shows the spectra of the local receive signal R_v in Fig. 5 when two sinusoids are remotely transmitted. The local transmitter is sending PAM8 data. The EC and QNC (when included) were initially adapted with the local transmitter on and the remote transmitter off; then the remote transmitter was turned on and an FFT was performed. The remote transmitted signal is not quantized; hence, the only noise source here is quantization noise from the DAC in the local transmitter. In this simulation, $B_1=12$ and $B_4=8$. The upper plot (thick line) is the spectrum without the QNC. The lower plot (thin line) is with the QNC. The two transmitted tones at 111 and 317 MHz are not affected by the QNC, but a significant reduction of the noise floor, which is due to DAC quantization noise, can be seen. The SNR (power in the sinusoids divided by power excluding the sinusoids) without the QNC is 41.3 dB, while the SNR with the QNC is 65.1 dB. The difference of 23.8 dB is very close to the expected 24.1 dB improvement if the QNC perfectly cancelled the 4-bit quantization noise.




Fig. 7. Spectra of the local R_v signal in Fig. 5 without QNC (thick line) and with QNC (thin line). Two sinusoids (111 and 317 MHz) are output by the remote transmitter. $B_1 = 12$ and $B_4 = 8$.

In the following simulations, the remote transmitter is off and the local transmitter is sending random PAM8 data, so the EC input (TxData) is a 3-bit word.

Fig. 8 shows two plots of MSE versus time with $B_1=12$ and $B_4=4$ (4-bit DAC). The mean-squared error, MSE (MSE = mean-squared value of err in Fig. 5, which includes uncancelled echo and uncancelled DAC quantization noise), is plotted with and without the QNC. The MSE is large initially and reduces as the EC and QNC (when included) converge. Without the QNC, the MSE converges to -18.1 dB. With the QNC, the MSE reaches -66.0 dB. Therefore, the reduction in MSE due the QNC is 47.9 dB, which is very close to the ideal improvement of 48.2 dB expected since the quantization noise $q[n]$ here is $B_1 - B_4 = 8$ bits.

An 8-bit DAC ($B_4=8$) was used to generate Fig. 9. The QNC reduces the MSE by 24.0 dB here, which is close to

the ideal reduction for 4-bit quantization noise. In this example, the input to the EC is 3-bit TxData, and the input to the QNC is 4-bit quantization noise $q[n]$.

The previous simulation was repeated with the same conditions ($B_1=12$, $B_4=8$) except a short, 16-tap QNC was used. This length QNC can cancel DAC quantization noise in the near echo but is not long enough to cancel DAC quantization noise in the far echo, which is much smaller than the near echo. The uncanceled quantization noise will appear in the R_v signal and increase the MSE. Fig. 10 plots the MSE vs. time with and without the short QNC. Here, the short QNC reduces the MSE by 16.8dB. In the previous simulation, the longer 128-tap QNC reduced the MSE by 24.0 dB. Therefore, the MSE increases by 7.2dB when the QNC length is reduced from 128 taps to 16 taps. In some systems, the MSE with the shorter QNC may be acceptable.

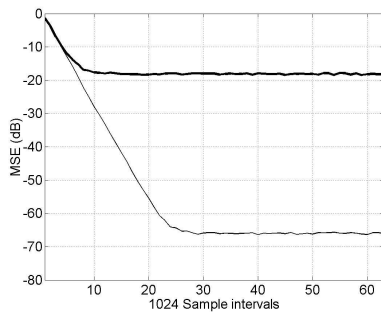


Fig. 8. Simulated MSE vs. time with 4-bit DAC and $B_1=12$, with QNC (thin line) and without QNC (thick line).




Fig. 9. Simulated MSE vs. time with 8-bit DAC and $B_1=12$, with QNC (thin line) and without QNC (thick line).

V. CONCLUSION

A quantization noise canceller has been proposed to cancel local transmit DAC quantization noise that appears in the local receiver in a full-duplex, baseband, echo-canceling transceiver. Simulations verify that the QNC effectively cancels the DAC quantization noise.

The QNC allows a reduced number of DAC bits. Fewer bits may lead to a DAC that is smaller, lower power, and/or more linear (due to better matching through the use of larger devices and passive components).

The simulations shown were carried out at a sampling rate of twice the baud rate. Similar QNC performance was observed when simulating with other sample rates.

The QNC can be used if the DAC input is a busy signal. A busy signal will exist if the PAM transmit data is significantly filtered digitally before the DAC, as may be the case when a digital low-pass or pre-emphasis filter operate on the (possibly upsampled) TxData.

The QNC can also be used with a modulated transmit signal. The key requirement is a linear signal path from the DAC input to the receiver signal from which the QNC output is subtracted. The QNC may be of use in full-duplex, echo-cancelled transceivers integrated in fine-line CMOS technologies where digital circuits are inexpensive and simple analog circuits are desired.




Fig. 10. Simulated MSE vs time with 8-bit DAC and $B_1=12$, with (thin line) and without (thick line) short 16-tap QNC.

REFERENCES

- [1] H. Samuelli, "Broadband Communications ICs: Enabling High-Bandwidth Connectivity in the Home and Office". *ISSCC Digest of Papers*, pp. 26-30 and *ISSCC Slide Supplement*, pp. 29-35, 1999.
- [2] K. Bult, "Analog Broadband Communication Circuits in Pure Digital Sub-Micron CMOS". *ISSCC*, pp. 76-77, 1999.
- [3] E. A. Lee and D. G. Messerschmitt, *Digital Communication*, Kluwer, New York, 1988.
- [4] J. A. C. Bingham, *The Theory and Practice of Modem Design*, Wiley, 1988.
- [5] A.V. Oppenheim and R.W. Schaffer, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, New Jersey, 1975, pp. 413-417.
- [6] A Gattani, D.W. Cline P.J. Hurst and P.M. Mosinskis, "A CMOS HDSL2 Analog Front-End", *IEEE J. of Solid-State Circuits*, pp. 1964-1975, Dec. 2000.
- [7] T.-C. Lee and B. Razavi, "A 125-MHz Mixed-Signal Echo Canceller for Gigabit Ethernet on Copper Wire", *IEEE J. of Solid-State Circuits*, pp. 366-373, Mar. 2001.
- [8] J. Cao, H. Lin, Y. Xiang, C. Kao, and K. Dyer, "A 10-bit 1GSample/s DAC in 90nm CMOS for Embedded Applications". *Custom IC Conf.*, pp. 165-168, 2006.
- [9] J. H. Fischer, J.L. Sonntag, J.S. Lavranchuk, D.P. Ciolini, A. Ganesan, D.G. Marsh, W.E. Keasler, J. Plany and L.H. Young, "Line and Receiver Interface Circuit for High-Speed Voice-Band Modems", *IEEE Journal of Solid State Circuits*, pp. 982-989, Dec. 1987.
- [10] J. Guido, V. Leung, J. Kenney, J. Trackim, A. Argillo, E. Zimany, and R. Shariadoust, "Analog Front End IC for Category I & II ADSL". *Symp. on VLSI Circuits*, pp. 178-181, 2000.
- [11] G. Zimmerman and B. Tolley, "10GBASE-T: 10G Ethernet on Copper Cable -- Part II," <http://www.networksystemsdesignline.com/howto/ethernet/showarticle.jhtml?articleID=184400885>
- [12] W. Y. Chen, *DSL Simulation Techniques and Standards Development for Digital Subscriber Line Systems*, Macmillan, Indianapolis, Indiana, 1998.
- [13] IEEE 802.3an standard, 10GBASE-T: 10Gb/s over copper.