

MIT Open Access Articles

*Exploiting cross-channel quantizer error correlation
in time-interleaved analog-to-digital converters*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: McMichael, Joseph G., Shay Maymon, and Alan V. Oppenheim. Exploiting Cross-channel Quantizer Error Correlation in Time-interleaved Analog-to-digital Converters. In 2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR), 525-529. Institute of Electrical and Electronics Engineers, 2011.

As Published: <http://dx.doi.org/10.1109/ACSSC.2011.6190056>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/79669>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Exploiting Cross-Channel Quantizer Error Correlation in Time-Interleaved Analog-to-Digital Converters

Joseph G. McMichael*, Shay Maymon†, and Alan V. Oppenheim
 Massachusetts Institute of Technology
 Digital Signal Processing Group
 77 Massachusetts Avenue, Cambridge MA 02139
 jmc Micha@mit.edu, maymon@mit.edu, avo@mit.edu

Abstract—Uniform quantizers are often modeled as additive uncorrelated noise sources. This paper explores the validity of the additive noise model in the environment of time-interleaved A/D converters. Cross-channel quantizer error correlation is an important discrepancy that arises for channel time delays in close proximity. It is demonstrated through simulation that negative error correlation occurs for different granularity quantizers in close proximity. Statistical analysis is presented to characterize error correlation between quantizers with different granularity. A technique exploiting this correlation often yields significant performance gains above the optimal additive noise model solution.

I. INTRODUCTION

TIME-INTERLEAVED analog-to-digital converters (ADCs) are often employed to sample high bandwidth signals or to achieve a large oversampling ratio. These devices interleave sample times of multiple converters operating at the same moderate rate. Traditionally, the relative time delays between channels have been selected to form an overall uniform sampling pattern, with equal quantization granularity in each channel. Recent work suggests that it is often possible to achieve a better signal-to-quantization noise ratio (SQNR) with different quantization granularity in each channel, non-uniform sampling, and appropriate reconstruction filtering ([1],[2],[3],[4]). This analysis relies on the additive noise model, which represents the effect of quantization error as an additive uncorrelated random process with uniform distribution that is also uncorrelated with the input [5]. In its traditional formulation, the additive noise model does not account for error correlation between channels.

In this paper, we explore the validity of the additive noise model in the environment of time-interleaved ADCs. We also present a technique that exploits cross-channel error correlation that often yields significant performance gains over the solutions based solely on the additive noise model.

This work was supported in part by a Fulbright Fellowship, Irwin and Joan Jacobs Presidential Fellowship, Texas Instruments Leadership University Program, BAE Systems PO 112991, Analog Devices, and Lincoln Laboratory PO 3077828.

*This author is now with MIT Lincoln Laboratory, Lexington, MA.

†This author is now with IBM T.J. Watson Research, Yorktown Heights, NY.

II. MULTI-CHANNEL SAMPLING AND RECONSTRUCTION

The basic multi-channel sampling and reconstruction system is shown in Fig. 1. In the measurement stage, the bandlimited signal $x(t)$ with Nyquist rate $1/T_N$ is uniformly sampled at the rate $1/(LT_N)$ in M parallel channels, where it is assumed that $M > L$. Although each channel undersamples $x(t)$ by a factor of L , the interleaved constellation has an effective oversampling ratio of $\rho = M/L > 1$. We denote the time delays in one recurrence period by τ_m and quantizer bit allocation by the $(1 \times M)$ vector \mathbf{b} , where the m^{th} entry corresponds to the number of bits in channel m .

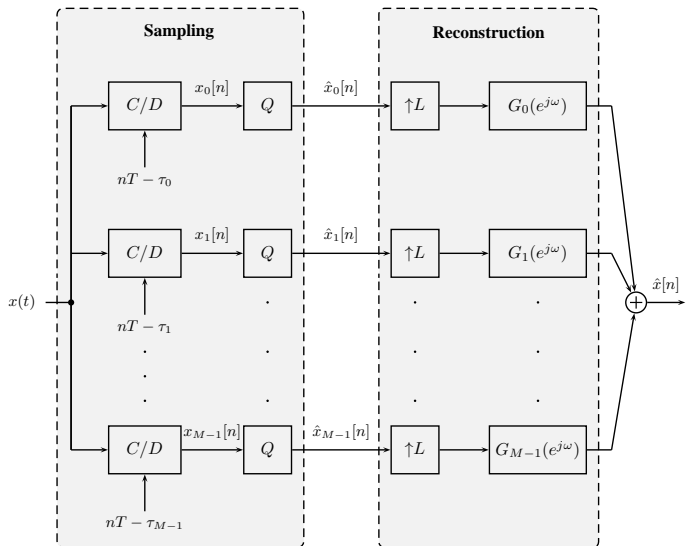


Fig. 1. Multi-channel sampling and reconstruction system.

When the time delays correspond to complex exponentials $e^{j2\pi\frac{\tau_m}{LT_N}}$ equally spaced on the complex unit circle, a uniform sampling grid is obtained. Otherwise, the interleaved grid can be characterized as recurrent non-uniform sampling ([6],[7]). The goal of the second stage in Fig. 1 is to compensate for non-uniform time delays and quantization error in the reconstruction.

A well-established analysis technique developed by Widrow and extended by others ([5],[8],[9],[10]) models quantization

error as additive uniformly distributed white noise. Specifically, the quantizer error in each channel is assumed to be a random process uniformly distributed between $\pm\Delta_m/2$ and uncorrelated with the input, where Δ_m denotes the quantizer step size in each channel. Correspondingly, the variance of the additive noise is given by:

$$\sigma_m^2 = \frac{\Delta_m^2}{12}. \quad (1)$$

The additive noise model performs quite well under many circumstances, particularly when the quantizer is not saturated, the granularity is small, and the probability density function (PDF) of the input is smooth ([9],[11]). In [1], [2], and [3], Maymon and Oppenheim derive optimal reconstruction filters that minimize the average noise power at the output of the reconstruction system under the additive noise model.

III. HOMOGENEOUS BIT ALLOCATION

Analysis in [1], [2], and [3] based on the additive noise model suggests that for a homogeneous bit allocation (i.e. with equal quantizer step size in each channel), the SQNR at the output of the reconstruction system in Fig. 1 is maximized when the time delays τ_m correspond to uniform sampling. To verify this result and explore any discrepancies, we have simulated the multi-channel system shown in Fig. 1. Consider the homogeneous system specified by $M = 3$, $L = 2$, $\mathbf{b} = [10, 10, 10]$, and $\tau_0 = 0$. We define the “reduction factor” γ in (2) as a measure of performance, where σ_m^2 denotes the error variance in the m^{th} channel and $\sigma_{e_{\min}}^2$ denotes error variance after optimal reconstruction filtering.

$$\gamma = \frac{\frac{1}{M} \sum_{m=0}^{M-1} \sigma_m^2}{\sigma_{e_{\min}}^2} \quad (2)$$

The simulation allows us to compare the effects of the additive noise model with those of true quantization. In the first scenario, a white uniformly distributed random process is added to each channel to model quantization error. Fig. 2 displays the reduction factor γ as a function of τ_1 and τ_2 with $\tau_0 = 0$, simulated with uniformly distributed error sources as specified by the additive noise model. Channel input signals are interpolated with recurrence frequency 4 kHz from a recording of the Norwegian Chamber Orchestra performing “Beethoven’s Symphony No. 1 in C,” bandlimited to 8 kHz with 24 bits per sample. As predicted, maximum noise reduction γ is achieved for $\frac{\tau_1}{T_N} = -\frac{\tau_2}{T_N} = \pm\frac{2}{3}$, which corresponds to uniform sampling.

The second scenario simulates a real system, where each channel includes a uniform quantizer. Fig. 3 displays the simulated reduction factor γ as a function of τ_1 and τ_2 , with $\tau_0 = 0$ and uniform quantizers. It is encouraging that the simulated performance based on the additive noise model matches the simulated performance with uniform quantizers well. The results reveal a large reduction in performance along the “groove” where the sampling instants of two channels fall in close proximity to each other. This reduction is not predicted by the analysis based on the additive noise model, which assumes channel error sources are uncorrelated. Samples of

a bandlimited signal approach identical values as their time separation approaches zero. Quantizing nearly identical values with the same granularity in a homogeneous system yields positively correlated error, significantly degrading performance after optimal reconstruction based on the uncorrelatedness assumption.




Fig. 2. Reduction factor γ as a function of delays τ_1 and τ_2 with the additive noise model for $\mathbf{b} = [10, 10, 10]$ and $\tau_0 = 0$. The two maxima of this surface correspond to uniform sampling patterns, affirming theoretical analysis.

Fig. 3. Reduction factor γ as a function of delays τ_1 and τ_2 with uniform quantizers for $\mathbf{b} = [10, 10, 10]$ and $\tau_0 = 0$. The results match those from Fig. 2 closely, except for three “grooves” with sampling instants in close proximity.

IV. HETEROGENEOUS BIT ALLOCATION

As we have seen, cross-channel quantizer error correlation arises for channel time delays in close proximity. This error correlation is positive for quantizers with the same granularity in close proximity, causing performance degradation. In this section, statistical analysis is presented showing that negative error correlation occurs for quantizers with different granularity in close proximity. With a diverse bit allocation, it is shown that error correlation can be exploited to improve performance.

The input-output relationship of a basic uniform, mid-tread quantizer is a staircase function, with step size Δ . For the purposes of this section, we will assume that the input does not saturate the quantizer. Increasing the number of bits by one doubles the number of levels and halves the maximum quantization error, as shown in Fig. 4.

For a quantizer with k bits and step size $\Delta = 2^{-(k-1)}$, the error signal ($e[n] = \hat{x}_m[n] - x_m[n]$) is “wrapped” to the interval $[-\frac{\Delta}{2}, \frac{\Delta}{2})$. We define E_k as the error for a k bit quantizer. The error for a $k+n$ bit quantizer can be obtained from the error E_k for a k bit quantizer by “wrapping” E_k n times, as shown in (3) for the special case of $n = 1$ and displayed in Fig. 5 for $n = 1, 2, 3$.

$$E_{k+1}(e_k) = \begin{cases} e_k + \Delta_{k+1} & \text{if } -\Delta_{k+1} \leq e_k < -\frac{\Delta_{k+1}}{2} \\ e_k & \text{if } -\frac{\Delta_{k+1}}{2} \leq e_k < \frac{\Delta_{k+1}}{2} \\ e_k - \Delta_{k+1} & \text{if } \frac{\Delta_{k+1}}{2} \leq e_k < \Delta_{k+1} \end{cases} \quad (3)$$

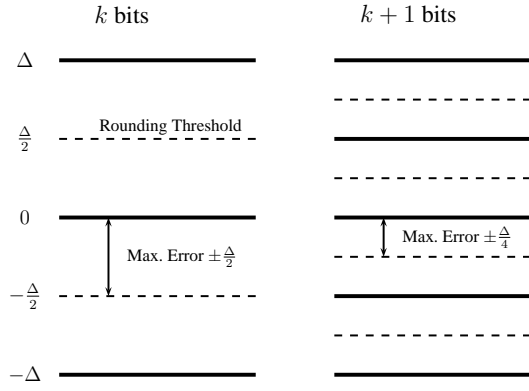


Fig. 4. Quantization regions and rounding thresholds (dashed lines) for k and $k + 1$ bits.

As shown in Fig. 5, more diagonal linear segments appear as the difference in bits n increases. The “base” region includes the linear segment centered around zero, from $-\frac{\Delta_{k+n}}{2}$ to $+\frac{\Delta_{k+n}}{2}$. The “middle” regions contain $2^{n-1} - 1$ shifted linear segments on each side, for $n \geq 2$. Finally, the two “outermost” regions extend to $\pm 2^{n-1} \Delta_{k+n}$ with length $\frac{\Delta_{k+n}}{2}$.




Fig. 5. Deterministic quantizer error functions for k bits to (a) $k + 1$ bits, (b) $k + 2$ bits, and (c) $k + 3$ bits.

A. Statistical Analysis of Quantizer Error Correlation

Samples of a bandlimited signal approach the same value as their time separation approaches zero. Correlation between channel inputs leads to correlation between the respective quantized error signals. The amount of correlation is dependent on the granularity of the two quantizers, and generally decays quickly with reduction in input correlation. Thus, for the remainder of this section we approximate channel inputs in close time proximity as identical analog values.

In this section, we analyze the statistical properties of the errors resulting from quantizing the same input sample with different granularities. Specifically, we consider the covariance and correlation coefficient of E_k and E_{k+n} . We begin with the definition of covariance in (4). We assume that E_k is uniformly distributed between $-\frac{\Delta_k}{2}$ and $\frac{\Delta_k}{2}$, as suggested by the additive

noise model. Thus,

$$\text{cov}(E_k, E_{k+n}) = \mathbb{E}[E_k E_{k+n}] - \mathbb{E}[E_k] \mathbb{E}[E_{k+n}] \quad (4)$$

$$= \mathbb{E}[E_k E_{k+n}] \quad (5)$$

$$= \int_{-\frac{\Delta_k}{2}}^{\frac{\Delta_k}{2}} e_k E_{k+n}(e_k) f_{E_k}(e_k) de_k. \quad (6)$$

Using our previous formulation of the function $E_{k+n}(e_k)$, we multiply by E_k to obtain the desired random variable $E_k E_{k+n}$. $E_k E_{k+n}$ contains quadratic segments over the same regions as $E_{k+n}(e_k)$. Fig. 6 depicts $E_k E_{k+3}$. The area under each of the quadratic segments has been shaded to emphasize that the expectation operation amounts to integration of this function scaled by a density constant for a uniform random variable.




Fig. 6. $E_k E_{k+3}$ as a function of E_k , where area has been shaded to emphasize that expectation reduces to integration.

Equation (6) can be simplified by noting that the PDF $f_{E_k}(e_k)$ is uniform, and thus evaluates to a constant.

$$\text{cov}(E_k, E_{k+n}) = \frac{1}{2^n \Delta_{k+n}} \int_{-2^{n-1} \Delta_{k+n}}^{2^{n-1} \Delta_{k+n}} e_k E_{k+n}(e_k) de_k \quad (7)$$

We separate the integral into the three regions, as shown in (8). Note from Fig. 6 that the product $E_{k+n}(e_k) E_k$ is an even function, allowing us to use symmetry to evaluate the integral more concisely.

$$\text{cov}(E_k, E_{k+n}) = \frac{1}{2^n \Delta_{k+n}} \left\{ \begin{array}{l} \text{base region} \\ \int_{-\frac{\Delta_{k+n}}{2}}^{\frac{\Delta_{k+n}}{2}} e_k^2 de_k \\ \text{middle regions} \\ + 2 \sum_{i=1}^{2^{n-1}-1} \int_{\Delta_{k+n}(i-\frac{1}{2})}^{\Delta_{k+n}(i+\frac{1}{2})} e_k^2 - i \Delta_{k+n} e_k de_k \\ \text{outermost regions} \\ + 2 \int_{(2^{n-1}-\frac{1}{2})\Delta_{k+n}}^{2^{n-1}\Delta_{k+n}} e_k^2 - 2^{n-1} \Delta_{k+n} e_k de_k \end{array} \right\} \quad (8)$$

Performing the integration in each region, we obtain (9). The base region and middle regions simply contribute the original positive variance of E_{k+n} weighted by their respective probabilities. However, the outermost region contributes a negative scaled version of the original variance. This occurs because the sign of E_k is always the opposite of E_{k+n} in

the outermost region. Combining terms, we arrive at (10) and (11), from which we conclude that the overall covariance $\text{cov}(E_k, E_{k+n})$ is simply negative one half of the variance of E_{k+n} .

$$\text{cov}(E_k, E_{k+n}) = \frac{\Delta_{k+n}^2}{12} \left(\overbrace{\frac{1}{2^n}}^{\text{base region}} + \overbrace{\left(1 - \frac{1}{2^{n-1}}\right)}^{\text{middle regions}} \right) + \overbrace{\left(1 - 3 \cdot 2^{n-1}\right) \left(\frac{1}{2^n}\right)}^{\text{outermost regions}} \quad (9)$$

$$= (-1/2) \frac{\Delta_{k+n}^2}{12} \quad (10)$$

$$= (-1/2) \text{var}(E_{k+n}) \quad (11)$$

By substituting the derived covariance and the variances of the respective uniform error signals, then using the relationship $\Delta_k = 2^n \Delta_{k+n}$, one may obtain the correlation coefficient ρ between E_k and E_{k+n} in (12).

$$\begin{aligned} \rho &= \frac{\left(-\frac{1}{2}\right) \frac{\Delta_{k+n}^2}{12}}{\sqrt{\left(\frac{\Delta_{k+n}^2}{12}\right) \left(\frac{\Delta_k^2}{12}\right)}} \\ &= \frac{\left(-\frac{1}{2}\right) \frac{\Delta_{k+n}^2}{12}}{\frac{\Delta_{k+n}^2 2^n}{12}} \\ &= -2^{-n-1} \end{aligned} \quad (12)$$

This reveals a simple formula for the error correlation between quantizers with the same input. As the difference in the number of bits n increases, the negative correlation is less pronounced. Equation (12) matches results from simulation well. Fig. 7 illustrates the predicted and measured correlation coefficient between a 1-20 bit quantizer output and a 10 bit quantizer output, with a uniformly distributed random input signal. The derived correlation measure no longer applies when $n = 0$ because the bit allocation is homogeneous.

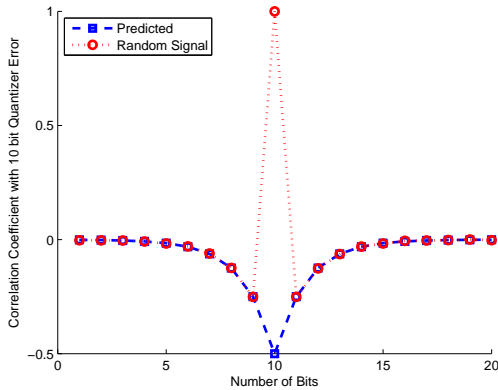


Fig. 7. Predicted and measured correlation coefficient between 1-20 bit quantizer error and 10 bit quantizer error.

B. Exploiting Correlated Error Signals

It is often possible to exploit error correlation to reduce quantization error. Various levels of design sophistication can be employed in a time-interleaved ADC to exploit this correlation. In fact, simply placing two quantizers at the same time instant with different granularities would clearly result in improvement in performance although this would not be predicted by the simple additive noise model.

Consider the system described earlier, with bit allocation $\mathbf{b} = [10, 10, 9]$ and $\tau_0 = 0$. Fig. 8 displays the reduction factor γ as a function of τ_1 and τ_2 , simulated based on the additive noise model. As discussed in [3], optimal constellations allocate high resolution quantizers greater time separation than low resolution quantizers. Fig. 9 displays the reduction factor γ when uniform quantizers are used. Interestingly, unpredicted improvements in performance are observed in the “groove” when the time separation between different granularity quantizers is small. In this case, negative cross-channel error correlation causes the two error signals to partially cancel during reconstruction, leading to performance improvements.




Fig. 8. Reduction factor γ as a function of delays τ_1 and τ_2 with the additive noise model for $\mathbf{b} = [10, 10, 9]$ and $\tau_0 = 0$.

Fig. 9. Reduction factor γ as a function of delays τ_1 and τ_2 with uniform quantizers for $\mathbf{b} = [10, 10, 9]$ and $\tau_0 = 0$. Unpredicted performance improvements occur when different granularity quantizers sample in close proximity.

More sophisticated methods can be devised to exploit the error correlation structure. A set of d quantized measurements \mathbf{x} in close proximity may be viewed as observations of the same analog input value μ combined with additive correlated error components with covariance matrix Σ . The Weighted Least Squares (WLS) criterion minimizes the sum of the weighted squares of the residuals to obtain an estimate for μ [12].

$$\mu_{WLS} = \underset{\mu}{\text{argmin}} (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu) \quad (13)$$

TABLE I
COVARIANCE MATRIX FOR ERROR IN $k - d + 1$ TO k BIT QUANTIZERS.

$$\Sigma = \begin{matrix} \text{bits} & k-d+1 & k-d+2 & \dots & k-1 & k \\ k-d+1 & +2^{2d-2} & -2^{2d-5} & \dots & -2^1 & -2^{-1} \\ k-d+2 & -2^{2d-5} & +2^{2d-4} & \dots & -2^1 & -2^{-1} \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ k-1 & -2^1 & -2^1 & \dots & +2^2 & -2^{-1} \\ k & -2^{-1} & -2^{-1} & \dots & -2^{-1} & +1 \end{matrix} \times \frac{\Delta_k^2}{12}$$

TABLE II
WLS WEIGHTS FOR d QUANTIZERS WITH BIT ALLOCATION $k - d + 1$ TO k .

$d \downarrow$	w_k	w_{k-1}	w_{k-2}	w_{k-3}	w_{k-4}
1	1	*	*	*	*
2	$\frac{3}{4}$	$\frac{1}{4}$	*	*	*
3	$\frac{2}{3}$	$\frac{1}{4}$	$\frac{1}{12}$	*	*
4	$\frac{5}{8}$	$\frac{1}{4}$	$\frac{3}{32}$	$\frac{1}{32}$	*
5	$\frac{3}{5}$	$\frac{1}{4}$	$\frac{1}{10}$	$\frac{3}{80}$	$\frac{1}{80}$

Taking the partial derivative with respect to μ and setting equal to zero yields the following weighted average, where $\mathbf{J} = \Sigma^{-1}$. The error covariance matrix Σ for an arbitrary bit allocation can be populated using (1) and (11).

$$\mu_{WLS} = \sum_{i=1}^d w_i x_i \quad \text{where} \quad w_i = \frac{\sum_{j=1}^d J_{i,j}}{\sum_{i=1}^d \sum_{j=1}^d J_{i,j}} \quad (14)$$

The covariance matrix for d quantizers with adjacent bit allocation $k - d + 1$ to k is displayed in Table I. Solving for w_i in (14) yields the set of weights in Table II, listed up to $d = 5$.

The theoretical effective resolution obtained through WLS unification of multiple quantizers was derived in [4]. For example, combining two quantizers ($d = 2$) with bit allocation $k - 1$ and k with the weights in Table II yields an effective resolution of $k + 0.339$ bits [4].

V. INCORPORATING CORRELATION INTO DESIGN

Finally, we incorporate the previous statistical analysis of cross-channel quantization error correlation into the design of optimal channel configurations. Initially, the channel configuration is optimized according to the additive noise model as outlined in [3]. For homogeneous bit allocations, imposing a minimum time separation between channels avoids performance degradation. For heterogeneous bit allocations, “unifying” groups of different-precision quantizers in close proximity using WLS can significantly improve performance. Optimal reconstruction filters are calculated with the effective resolution of the WLS unified quantizer derived in [4].

Consider the case when $M = 9$, $L = 6$, and $\mathbf{b} = [10, 10, 10, 9, 9, 9, 8, 8, 8]$. The optimal constellation based on the additive noise model is depicted in Fig. 10, where the circumference of the circle is equal to the recurrent sampling period $L \cdot T_N$. To obtain the optimal configuration for actual quantizers, we must “unify” the 8 and 9 bit quantizers in close proximity. Fig. 11 displays the simulated reduction factor γ with uniform quantizers for various levels of design sophistication. When the time delays are constrained to be uniform, the best possible reduction factor is $\gamma = 2.65$. Allowing non-uniform time delays without incorporating knowledge

of correlation boosts performance significantly to $\gamma = 3.31$. Correcting the 8 and 9 bit channel time delays so that each pair has the same sampling instant yields another significant performance improvement to $\gamma = 3.91$. Finally, keeping the previous corrections and performing WLS estimation yields another small improvement, to $\gamma = 3.98$. This corresponds to a 16.8% reduction in output error variance versus the predicted optimal time configuration, and a 33.4% reduction in output error variance versus the uniform sampling pattern. In general, systems with many channels and diverse bit allocations allow for larger performance improvements.




Fig. 10. Predicted optimal time delays for $M = 9$ and $L = 6$ found through numerical optimization, for $\mathbf{b} = [10, 10, 10, 9, 9, 9, 8, 8, 8]$. Superscripts denote the associated bit allocation for each channel.




Fig. 11. Best reduction factor γ for $\mathbf{b} = [10, 10, 10, 9, 9, 9, 8, 8, 8]$, with increasing levels of design sophistication.

REFERENCES

- [1] S. Maymon. *Sampling and Quantization for Optimal Reconstruction*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [2] S. Maymon and A.V. Oppenheim. Quantization and compensation in sampled interleaved multi-channel systems. Accepted for publication by IEEE Transactions on Signal Processing, Jan. 2012.
- [3] S. Maymon and A.V. Oppenheim. Quantization and compensation in sampled interleaved multi-channel systems. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 3622–3625. IEEE, 2010.
- [4] J.G. McMichael. Timing Offset and Quantization Error Trade-off in Interleaved Multi-Channel Measurements. Master’s thesis, Massachusetts Institute of Technology, 2011.
- [5] B. Widrow and I. Kollár. *Quantization noise: roundoff error in digital computation, signal processing, control, and communications*. Cambridge University Press, 2008.
- [6] A.J. Jerri. The Shannon sampling theorem - Its various extensions and applications: A tutorial review. *Proceedings of the IEEE*, 65(11):1565–1596, 2005.
- [7] J. Yen. On nonuniform sampling of bandwidth-limited signals. *Circuit Theory, IRE Transactions on*, 3(4):251–257, 1956.
- [8] A. Sripad and D. Snyder. A necessary and sufficient condition for quantization errors to be uniform and white. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 25(5):442–448, 2003.
- [9] W.R. Bennett. Spectra of quantized signals. *Bell Syst. Tech. J*, 27(3):446–472, 1948.
- [10] R.M. Gray. Quantization noise spectra. *Information Theory, IEEE Transactions on*, 36(6):1220–1244, 2002.
- [11] A.V. Oppenheim and R.W. Schaffer. *Discrete-time signal processing*. Prentice-Hall signal processing series. Prentice Hall, 2010.
- [12] S.M. Kay. *Fundamentals of Statistical Signal Processing: Estimation theory*. Prentice Hall signal processing series. Prentice-Hall PTR, 1998.