This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2010 proceedings

Channel Estimation with Low-Precision Analog-to-Digital Conversion

Onkar Dabeer
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai
India
Email: onkar@tcs.tifr.res.in

Upamanyu Madhow
Electrical and Computer Engineering
University of California
Santa Barbara
USA
Email: madhow@ece.ucsb.edu

***Abstract*— We consider the problem of estimating the impulse response of a dispersive channel when the channel output is sampled using a low-precision analog-to-digital converter (ADC). While traditional channel estimation techniques require about 6 bits of ADC precision to approach full-precision performance, we are motivated by applications to multiGigabit communication, where we may be forced to use much lower precision (e.g., 1-3 bits) due to considerations of cost, power, and technological feasibility. We show that, even with such low ADC precision, it is possible to attain near full-precision performance using closed-loop estimation, where the ADC input is dithered and scaled. The dither signal is obtained using linear feedback based on the Minimum Mean Squared Error (MMSE) criterion. The dither feedback coefficients and the scaling gains are computed offline using Monte Carlo simulations based on a statistical model for the channel taps, and are found to work well over wide range of channel variations.**

***Index Terms*— Channel estimation, multiGigabit communications, 60 GHz, analog-to-digital conversion**

I. Introduction

The analog-to-digital converter (ADC) is a key component in modern digital communication receivers allowing exploitation of Moore's law for low-cost implementation of sophisticated receiver functionalities in digital signal processing (DSP). As communication speeds scale up, however, the cost, power and availability of ADCs with sufficient precision (typical receiver implementations use 6-12 bits of precision) becomes a challenge. For example, there is significant activity in standardization of multiGigabit wireless personal area networks (IEEE 802.25.3c [6]) and wireless local area networks (IEEE 802.11ad [5]) using the unlicensed 60 GHz band. DSP-centric transceiver implementations in such settings would require ADCs with sampling rates of the order of several GHz. The power consumption at such speeds could be significantly reduced by drastically lowering ADC precision (e.g., to 1-3 bits), and recent information-theoretic results [8], [9], [10] show that, under ideal conditions (nondispersive channel, perfect synchronization), the capacity loss due to low-precision ADC is acceptable even at moderately high signal-to-noise ratio (SNR). This motivates more detailed investigation into whether advanced signal processing functionalities such as channel estimation and synchronization can be accomplished using DSP on low-precision samples. In this paper, our focus is on estimation of the impulse response of a dispersive channel when the noisy channel output is quantized using low-precision ADCs.

The work of Onkar Dabeer was supported by TIFR under the XI Plan Project.

The work of Upamanyu Madhow was supported by the National Science Foundation under grants CCF-0729222 and ECS-0636621.

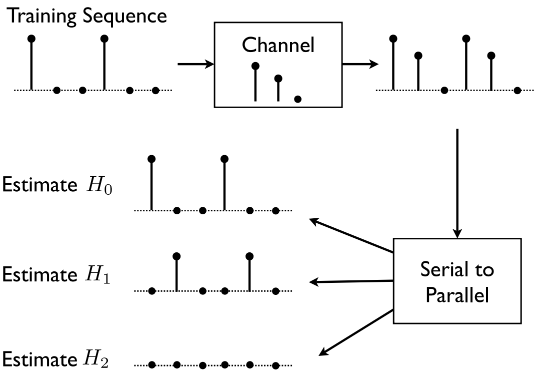

Fig. 1. Reduction of channel estimation to $K = 3$ mean estimation problems using a periodic bursty training sequence.

A. Problem Definition

Consider an inter-symbol interference channel with impulse response of length *at most* K and denote the taps by $\{H_0, ..., H_{K-1}\}$. If the input symbols to the channel are denoted by $\{X_t\}$, then the noiseless full-precision output is given by $Z_t = \sum_{k=0}^{K-1} H_k X_{t-k}$. If we use a periodic training sequence (i.e., choose $\{X_t\}$ to be periodic), then $\{Z_t\}$ is also periodic. In this case, channel estimation can be decomposed into a number of parallel mean estimation problems. This is easiest to see for a bursty training sequence with period K, which during one period sends a nonzero training symbol followed by $K-1$ zeros. In this case, a single period of $\{Z_t\}$ simply equals the channel impulse response (padded with zeros if the response is shorter than K). As illustrated in Figure

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2010 proceedings

1, this allows us to decompose the channel estimation problem into K parallel single parameter estimation problems: Estimate H_k based on the channel output at times $\{k + nK, n \geq 0\}$. However, note that such a decomposition applies to a standard pseudorandom training sequence as well: as long as it is periodic with period K, so is $\{Z_t\}$, which implies that we can focus on estimating, say, Z_k based on observations at times $k, k + K, k + 2K, \ldots$. Once we have accurate estimates of $\{Z_t\}$, we can estimate the channel coefficients as usual by deconvolution with the known training sequence.

The preceding decomposition allows us to focus on the fundamental problem of estimating $A = Z_k$, based on the noisy quantized observations $Y_n = \mathsf{Q}_n(A + W_n)$, $1 \leq n \leq N$ (these observations are spaced by the period K of the training sequence), where $\{W_n\}$ is complex white Gaussian noise and Q_n denotes the quantizer used for sample n. Our goal is to design the quantizers $\{\mathsf{Q}_n\}$ so as to get a good estimate for A with the smallest number, N, of samples. For a training sequence of period K, we have K such estimators running in parallel. Thus, if N observations are needed for each of these estimators, then the required training sequence length is KN (ignoring edge effects).

B. Summary of Results

Having reduced the channel estimation problem to a number of parallel mean estimation problems, we show that closed loop estimation with a dither signal based on linear feedback provides excellent performance. As in sigma-delta converters [1], the aim of the closed loop dither signal is to estimate the input and reduce the dynamic range of the signal entering the quantizer. Specifically, we choose the dither signal to be the linear MMSE estimate of the desired coefficient based on the quantized observations so far, and the feedback taps are designed offline using a Gaussian prior on A. In addition, we apply an open loop time-varying gain (designed offline as well) after dithering but before quantization. This is in addition to the automatic gain control (AGC) *prior to* dithering being used in the receiver, and is designed to match the dynamic range of the signal *after dithering* to that of the quantizer. Our results show that the dither signal and the gain signal substantially improve performance relative to an open loop approach, and provide performance close to that with full-precision sampling. For SNR = 10 dB, our estimator based on 3-bit (2-bit) ADC gives a MSE of -20 dB with 11 (13 respectively) training symbols, compared with 10 training symbols for full-precision sampling. Even though we assume a Gaussian prior, we find that the performance of our estimator is quite insensitive to the actual channel realization.

C. Related Work

The problem of signal parameter estimation based on 1-bit samples has received much attention in recent years; see for example [2], [3], [7] and references therein. These works bring out the critical role of dithering to combat the severe non-linearity of low-precision quantization. For example, in [7] it is shown that for the mean estimation problem, adding dither before quantization can improve the Cramer-Rao lower bound (CRLB), and it is shown in [2], [3] that for a general class of signal parameter estimation problem, the least-squares fit estimator can be made consistent by using a dither signal. Our work differs from the preceding papers in several respects. Since we consider a general finite precision ADC rather than 1-bit ADC, we introduce gain control in addition to dither. While [7] optimizes the dither signal based on the Cramer-Rao Lower Bound (CRLB), we are interested in low to moderate SNR and short training sequences, so that the CRLB (which applies in the asymptotic regimes corresponding to a large number of samples or high SNR) is less relevant. Instead, we choose a linear MMSE criterion for design of our dither signal.

D. Organization

The paper is organized as follows. In Section II, we state our model and assumptions. In Section III-A, we consider closed loop design of the dither signal using time-varying taps; in Section III-B, we consider time-invariant feedback taps; in Section III-C, we consider design of the open loop gain signal; and in Section III-D, we show that our design is insensitive to the Gaussian prior assumed on the channel. The conclusion is given in Section IV.

II. System Model

Consider

$$Y_n = \mathsf{Q}_n(A + W_n), \quad 1 \leq n \leq N,$$

where $\{W_n\}$ are i.i.d. $\mathbb{CN}(0, \sigma^2)$, A is $\mathbb{CN}(0, \sigma_A^2)$, and the quantizer sequence Q_n is to be designed based on a finite-precision ADC with a small number (e.g., 1-3 bits) of bits of precision. In standard receivers in which the available ADC precision is large enough (e.g., 6-12 bits), the quantizer is typically chosen to be time-invariant and its presence is ignored (i.e., estimation algorithms are designed assuming that the available samples are at full precision). However, once we restrict the ADC to, say, 1-3 bits of precision, the quantizer nonlinearity is severe and cannot be ignored.

Quantizer Structure: From prior work on 1-bit ADC [2], [3], [7], we know that dithering can significantly improve performance for such nonlinear estimation problems. We therefore consider quantizers of the following form:

$$\mathsf{Q}_n(x) = \mathsf{Q}\left(G_n(x - V_n)\right) \tag{1}$$

where $\{V_n\}$ is a dither signal to be designed, $\{G_n\}$ is a positive gain signal to be designed, and $\mathsf{Q}(\cdot)$ is the uniform quantizer of precision 2ℓ bits per complex-valued sample (ℓ bits for real part and ℓ bits for imaginary part)

$$\begin{aligned}\mathsf{Q}(x) &= \mathsf{Q}_u(\mathrm{Re}(x)) + i\mathsf{Q}_u(\mathrm{Im}(x)) \\ \mathsf{Q}_u(u) &= \mathrm{sign}(u)\left[\delta\left\lfloor\frac{|u|}{\delta}\right\rfloor + \frac{\delta}{2}\right], \quad |u| \leq 2^{\ell-1}\delta \\ &= \mathrm{sign}(u)\frac{(2^\ell - 1)}{2}\delta, \text{ otherwise.}\end{aligned}$$

While all our techniques are valid for any finite ℓ, our primary interest is in 1-3 bit ADCs ($\ell = 1, 2, 3$).

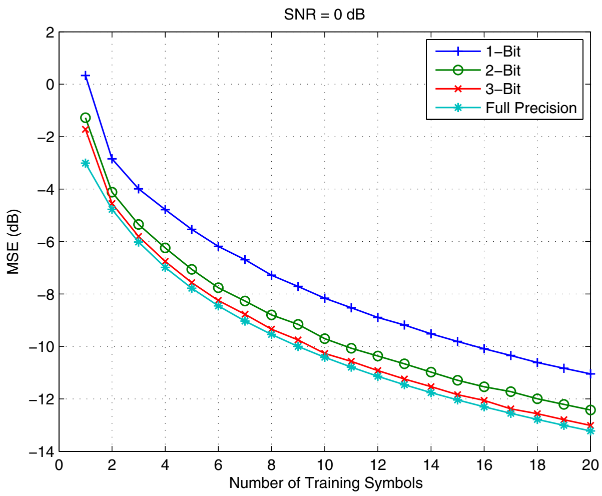

Fig. 2. 3-bit LMMSE CL-ADC is close to full precision MMSE estimate for 0 dB SNR. For a target MSE of -10 dB, $N = 9$ for full precision, while 1-bit, 2-bit, and 3-bit LMMSE CL-ADCs need $N = 16$, $N = 11$, $N = 10$ respectively.

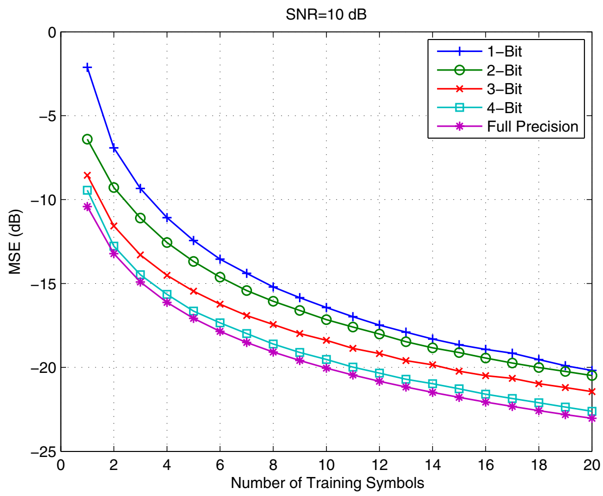

Fig. 3. 4-bit LMMSE CL-ADC is close to full precision MMSE estimate for 10 dB SNR. For a target MSE of -20 dB, $N = 10$ for full precision, while 1-bit, 2-bit, and 3-bit LMMSE CL-ADCs need $N = 19$, $N = 18$, $N = 14$ respectively.

Statistical Prior: Our goal in this paper is to design simple schemes to adapt the quantizer parameters (V_n, G_n). To this end, we assume that A is complex Gaussian with zero mean and known variance. For a bursty training sequence, this amounts to assuming that the channel is Rayleigh fading. For a standard training sequence, the noiseless channel output at any time is a weighted sum of a number of elements of the training sequence, where the weights depend on the channel impulse response. If the channel impulse response contains terms of comparable magnitude, then the central limit theorem suggests a Gaussian prior. Even though we assume a specific prior, we find (see Section III-D) that the performance of our estimator is insensitive to the actual realization of A.

The signals $\{V_n, G_n\}$ allow us to look at the mean A through a family of quantizers obtained from $\mathsf{Q}(\cdot)$. These can be chosen in two ways.

1) **Open Loop ADC (OL-ADC):** In this case, $\{V_n, G_n\}$ are designed offline and do not depend on the realization of the data. To gain insight, it is worthwhile to look at the simple case when $\ell = 1$. In this case, the gain signal does not play any role and we set $G_n = 1$. Further, if we let $\sigma \to 0$, then the samples Y_1^N correspond to N comparisons, and we get a resolution of $O(1/N)$.
2) **Closed Loop ADC (CL-ADC):** In this case, V_n, G_n are chosen online depending on the past samples Y_1^{n-1}. If we again consider $\ell = 1$, then in the limiting case $\sigma \to 0$, we can get a resolution of $O(1/2^N)$, which is exponentially better than the OL-ADC. In [7], it is also shown that the CL-ADC has a substantially better CRLB than the OL-ADC.

To simplify implementation of the CL-ADC, we make two choices:

- we adapt V_n using feedback, but we simplify the feedback structure by considering only linear combinations of past observations;
- we design G_n offline.

Moreover, our final estimate of A also depends linearly on Y_1^N. In [7], the case $\ell = 1$ (and hence $G_n = 1$) is considered and numerical optimization of the CRLB is used to propose good dither signals and mean estimators. Since we are interested in small N and simple linear estimators of A given Y_1^N, the CRLB may not be relevant to our case. Instead, we rely on the minimum mean-square error (MMSE) criterion.

III. Main Results

In this section, we first consider linear feedback dither signal with time-varying taps and constant gain ADC. This is followed by evaluation of impact of time-invariant feedback taps, time-varying gain signal, and performance for a fixed channel realization.

A. Linear Feedback Dither and Constant Gain ADC

In this section, we consider $G_n = 1$ and $V_n = \sum_{t=1}^{n-1} b_{n-1,t}^* Y_t = \boldsymbol{b}_{n-1}^H \boldsymbol{Y}_{n-1}$ where $\boldsymbol{Y}_n = [Y_1, ..., Y_n]^T$ and $\boldsymbol{b}_n = [b_{n,1}, ..., b_{n,n}]^T$. We design V_n to be a linear MMSE estimate of A, so that at a given step we are quantizing the error with respect to the estimate based on the observations in previous steps. Since $\boldsymbol{Y}_0$ is empty, we take $V_1 = 0$. While we allow time-varying feedback taps, we show later that appropriately chosen time-invariant taps work almost as well. The linear MMSE (LMMSE) estimator at time n is

$$\boldsymbol{b}_n^{mmse} = \mathbf{R}^{-1}(n)\boldsymbol{r}(n),$$

where

$$\boldsymbol{r}(n) = E[A^* \boldsymbol{Y}_n], \quad \mathbf{R}(n) = E[\boldsymbol{Y}_n \boldsymbol{Y}_n^H].$$

We refer to this CL-ADC as the LMMSE CL-ADC. We note that $\boldsymbol{r}(n)$, $\mathbf{R}(n)$ depends on $\boldsymbol{b}_1, ..., \boldsymbol{b}_{n-1}$. Thus we can compute the feedback weights $\boldsymbol{b}_n^{mmse}$ by using the following recursion.

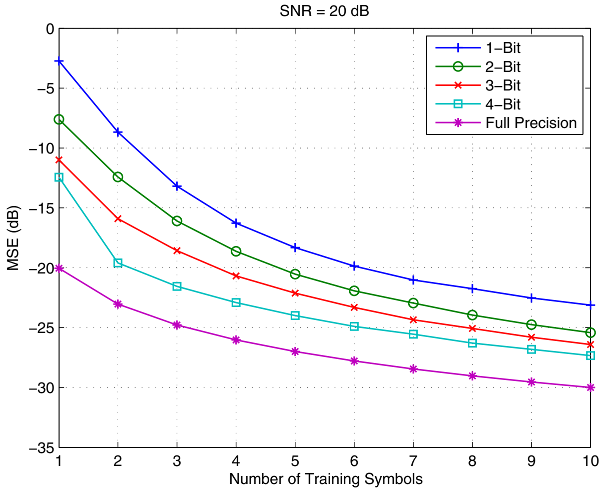

Fig. 4. For SNR of 20 dB, 1-4 bit ADCs do not approach full precision performance, but still yield excellent performance. For a target MSE of -20 dB, $N = 1$ for full precision, while 1-bit, 2-bit, 3-bit and 4-bit LMMSE CL-ADCs need $N = 6$, $N = 5$, $N = 4$, $N = 2$ respectively.

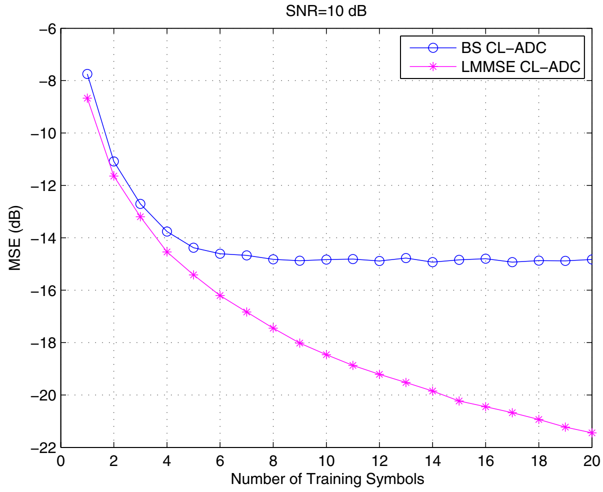

Fig. 5. LMMSE CL-ADC significantly outperforms BS CL-ADC when the target MSE is small.

1) Start with $n = 1$.
2) Compute $\mathbf{R}(n)$ and $\boldsymbol{r}(n)$.
3) Compute $\boldsymbol{b}_n^{mmse} = \mathbf{R}^{-1}(n)\boldsymbol{r}(n)$.
4) Increment n by 1 and go to step 2.

The optimal feedback weights are computed based on nominal values of σ_A^2, σ^2. Unfortunately, due to the dependence caused by the feedback and the non-linearity of the quantizer, in general there are no closed form expressions for Step 2) above. Hence we resort to Monte-Carlo simulations (over the statistics of A and the noise) to estimate $\mathbf{R}(n)$ and $\boldsymbol{r}(n)$. The MSE of the resulting estimators for 1-4 bits of ADC precision is plotted for different SNR in Figures 2, 3, 4. As a benchmark, we also plot the performance of the full precision MMSE estimator:

$$\text{MSE}_* = \frac{\sigma^2}{N\sigma_A^2 + \sigma^2}.$$

Without loss of generality, we normalize $\sigma_A^2 = 1$. Since the real and imaginary parts of A lie in $[-3/\sqrt{2}, 3/\sqrt{2}]$ with high probability, we set the quantizer dynamic range to

$$2^{\ell-1}\delta = \frac{3}{\sqrt{2}}.$$

to match the range of A. We find that 3-bit ADCs ($\ell = 3$) yield excellent performance:

- for SNR of 0 dB, it needs $N = 10$ for a MSE of -10 dB compared to $N = 9$ for full precision;
- for SNR of 10 dB, it needs $N = 14$ for a MSE of -20 dB compared to $N = 10$ for full precision;
- for SNR of 20 dB, it needs $N = 4$ for a MSE of -20 dB compared to $N = 1$ for full precision.

In Section III-C, we show that this performance can be improved further by using gain adaptation.

B. Time Invariant Feedback Taps

To simplify implementation, we are interested in time invariant feedback tap settings. We compare a benchmark "binary search" approach with taps obtained using the LMMSE criterion. We consider two cases.

1) **"Binary search" CL-ADC:** The taps for this benchmark are given by

$$b_{n,k} = \frac{\Delta}{2^k}.$$

For 1-bit quantization and $\sigma = 0$, these weights can be interpreted as a binary search for the value of A based on 1-bit samples. Hence we refer to this ADC as the binary search (BS) CL-ADC. For SNR = 10 dB, in Figure 5 we compare the performance of the LMMSE CL-ADC with BS CL-ADC. If the target MSE is larger than -14 dB, then we see that LMMSE and BS CL-ADC are close. But the MSE of the BS CL-ADC saturates at around -15 dB while the MSE of LMMSE CL-ADC continues to decrease with N.

2) **Time invariant LMMSE:** As we iterate over N in our simulations to find $\boldsymbol{b}_n^{mmse}$, we find that after the first few iterations $\boldsymbol{b}_{n-1}^{mmse}$ is very close to the first $n-1$ entries of $\boldsymbol{b}_n^{mmse}$. For example, for $\ell = 3$, SNR=10 dB, the feedback tap matrix $[b_{i,j}]$ (with row i as $\boldsymbol{b}_i^t$) for size $n = 5$ is:

$$\begin{bmatrix} 0.7979 & 0 & 0 & 0 & 0 \\ 0.8401 & 0.4705 & 0 & 0 & 0 \\ 0.8405 & 0.4673 & 0.3182 & 0 & 0 \\ 0.8407 & 0.4693 & 0.3166 & 0.2427 & 0 \\ 0.8399 & 0.4666 & 0.3177 & 0.2429 & 0.1919 \end{bmatrix}.$$

This suggests that we can use $\boldsymbol{b}_N^{mmse}$ as the time invariant filter taps, that is,

$$b_{n,k} = b_{N,k}, \quad k \leq n.$$

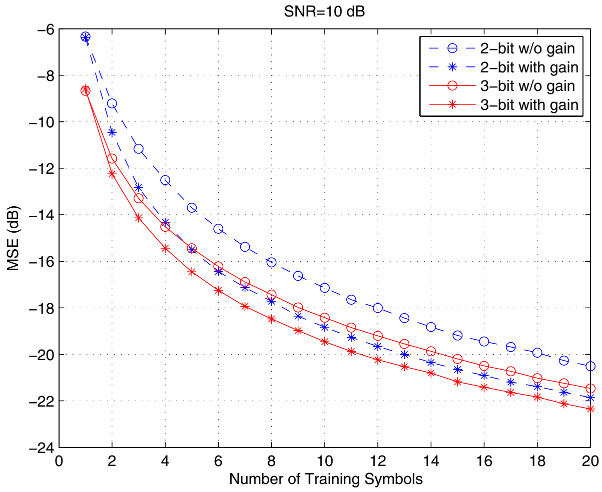

Fig. 6. At SNR=10 dB, 2-bit ADC with gain control outperforms 3-bit ADC with constant gain for $N > 5$. Moreover 2-bit ADC with gain control needs $N = 13$ to get MSE of -20 dB compared to $N = 18$ without gain control and $N = 10$ for full precision case, while 3-bit CL-ADC with gain control needs only $N = 11$.

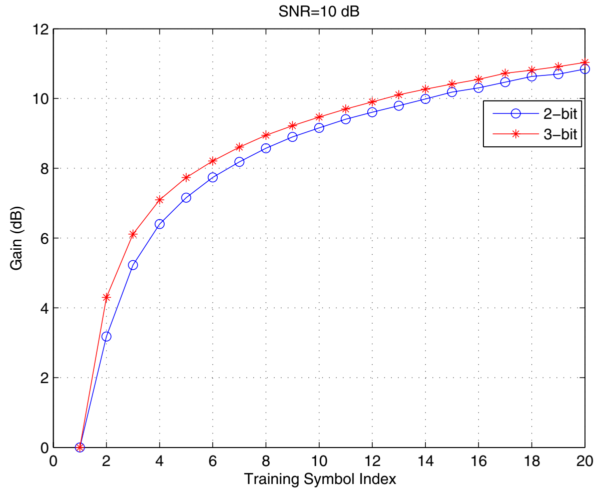

Fig. 7. The gain signal (designed offline) scales up the error as the number of observations increases, in order to fully utilize the ADC dynamic range.

We find that this feedback filter gives almost identical performance to LMMSE CL-ADC.

C. Offline Design of Gain Signal

To understand the limitation of using a constant gain ($G_n = 1$), consider a situation in which the dither V_n is a good estimate of the signal A so that the dynamic range of $A - V_n$ is smaller than the quantizer resolution. In this case, even when $\sigma = 0$, the multilevel quantizer essentially gives only 1 bit of information, so that scaling the error $A - V_n$ to match the dynamic range of the quantizer can be expected to provide performance gains. At the nth sample, we quantize $A - V_n$, and $\sigma_n^2 = E[|A - V_n|^2]$ is a measure of the dynamic range of the signal. If $A - V_n$ is Gaussian, then with high probability its real and imaginary parts lie in the interval

$$\left[-3\sigma_n/\sqrt{2}, 3\sigma_n/\sqrt{2}\right].$$

The quantizer dynamic range is $[-2^{\ell-1}\delta, 2^{\ell-1}\delta]$. Assuming that the Gaussian approximation is good, we can match the dynamic ranges by choosing

$$G_n = \frac{\sqrt{2}2^{\ell-1}\delta}{3\sigma_n}.$$

In Figure 6, we study the impact of such gain control for SNR=10 dB. We see that gain control greatly improves the 2-bit LMMSE CL-ADC: with gain control we need $N = 13$ to get MSE of -20 dB compared to $N = 18$ without gain control and $N = 10$ for full precision case. Moreover, the 2-bit LMMSE CL-ADC with gain control is better than the 3-bit LMMSE CL-ADC without gain control for $N > 5$. The 3-bit LMMSE CL-ADC with gain control needs $N = 11$ to attain a MSE of -20 dB compared with $N = 10$ for full precision. In Figure 7, we plot the values of G_n.

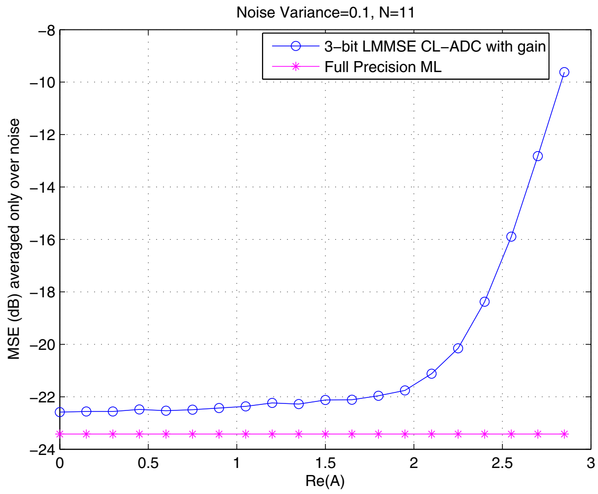

Fig. 8. Our estimator is insensitive to the actual realization of A within the quantizer dynamic range [-2.12,2.12].

For SNR = 0 dB, 2-bit and 3-bit LMMSE CL-ADC are close to full precision, and not surprisingly, gain control has little impact. For SNR = 20 dB, 2-bit and 3-bit LMMSE CL-ADC with gain control need $N = 3$ and $N = 4$ respectively to achieve a MSE of -20 dB compared to $N = 5$ and $N = 4$ respectively in the absence of gain control. Thus, the proposed open loop gain control can significantly reduce the training overhead in moderate to high SNR regime.

D. MSE for a Given Realization of A

In the previous sections, the MSE was averaged over the noise in the data as well as the prior on A. In this section, we consider the MSE averaged only over the noise to understand the performance for a fixed value of A:

$$\mathcal{E}(A) = E[(A - V_N)^2|A].$$

In Figure 8, we plot

$$\mathcal{E}'(\mathrm{Re}(A)) = E[(\mathrm{Re}(A - V_N))^2 | \mathrm{Re}(A)]$$

for the case of full precision sample mean estimator and 3-bit LMMSE CL-ADC with gain control for $N = 11$. The quantizer dynamic range is chosen to be $[-3/\sqrt{2}, 3/\sqrt{2}] \approx [-2.12, 2.12]$, and since A is zero mean Gaussian with unit variance, with very high probability $\mathrm{Re}(A)$ belongs to this range. Due to symmetry, in Figure 8, we only plot for the positive values of $\mathrm{Re}(A)$. We see that for $\mathrm{Re}(A)$ inside the nominal quantizer dynamic range, $\mathcal{E}'(\mathrm{Re}(A))$ is quite insensitive to $\mathrm{Re}(A)$, indicating that quantizer design and estimator based on the Gaussian prior is robust. However, once the quantizer dynamic range is exceeded, information about the signal is lost, and as expected, the performance degrades rapidly.

IV. Conclusion

We have shown that mean estimation can be accomplished effectively with low-precision ADC using linear MMSE feedback and gain control. However, there are a large number of issues that remain to be investigated as to how such mean estimation can be used as a building block for communication transceiver design. The K parallel mean estimation problems considered yield estimates of the noiseless channel output, and directly provide estimates of the channel impulse response for a bursty training sequence. However, for a standard training sequence, we must deconvolve the training sequence to obtain estimates of the impulse response; while we do not expect any surprises, it is important to evaluate the sensitivity of such a deconvolution to errors in the parallel mean estimates. Once the channel has been estimated, we have two basic choices: we can use the channel estimates to design equalizers at the receiver (taking into account the nonlinearity caused by the low precision of the samples) and/or feed back the channel estimates to the transmitter, so that it can use precoding ([4]) to compensate for the channel. The latter is especially attractive for asymmetric links with a more powerful transmitter (e.g., a laptop transmitting to a handheld on a 60 GHz link).

The proposed feedback from the digital to analog part of the receiver requires a digital-to-analog converter (DAC). In future, we plan to analyze the impact of the placement and accuracy of the DAC on complexity and performance of the receiver.

Finally, we must address the joint design of carrier synchronization and channel dispersion compensation for receivers with low-precision ADC (see [11] for a Shannon theoretic analysis of carrier synchronization for a *non-dispersive* channel).

References

[1] P. M. Aziz and H. V. Sorensen and Jan Van der Spiegel, “An Overview of Sigma-Delta Converters: How a 1-bit ADC achieves more than 16-bit resolution,” IEEE Signal Processing Magazine, Volume 13, Issue 1, September 1996, pages 61-84.

[2] O. Dabeer and E. Masry, “Multivariate Signal Parameter Estimation Under Dependent Noise From 1-bit Dithered Quantized Data,” *IEEE Transactions on Information Theory*, vol. 54, no. 4, pp. 1637-1654, April 2008.

[3] O. Dabeer and A. Karnik, “Signal Parameter Estimation with Dithered 1-bit Quantization,” *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5389-5405, December 2006.

[4] Harashima, H.; Miyakawa, H., “Matched-Transmission Technique for Channels With Intersymbol Interference,” *IEEE Transactions on Communications*, Aug 1972 Volume: 20, Issue: 4, page(s): 774- 780.

[5] IEEE 802.11 Task Group AD, http://www.ieee802.org/11/Reports /tgadupdate.htm

[6] IEEE 802.15.3c Task Group, http://www.ieee802.org/15/pub/TG3c.html

[7] H. C. Papadopoulos, G. W. Wornell, and A. V. Oppenheim, “Sequential signal encoding from noisy measurements using quantizers with dynamic bias control,” *IEEE Transactions on Information Theory*, vol. 47, no. 3, pp. 978–1002, March 2001.

[8] J. Singh, O. Dabeer, U. Madhow, “Transceiver Design with Low-Precision Analog-to-Digital Conversion : An Information-Theoretic Perspective,” to appear *IEEE Transactions on Communications*, 2009.

[9] J. Singh, O. Dabeer and U. Madhow, “Capacity of the Discrete-Time AWGN Channel Under Output Quantization,” *IEEE International Symposium on Information Theory*, July 2008, Toronto, Canada.

[10] J. Singh, O. Dabeer, U. Madhow, “Communication Limits with Low Precision Analog-to-Digital Conversion at the Receiver,” Proceedings of the *IEEE International Conference on Communications*, pp. 6269-6274, Glasgow, June 2007.

[11] J. Singh and U. Madhow, “On block noncoherent communication with low-precision phase quantization at the receiver,” *Proc. 2009 International Symposium on Information Theory (ISIT 2009)*, South Korea, July 2009.

[12] R. Walden, “Analog-to-Digital Converter Survey and Analysis,” *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, April 1999.