Space-time Slicer Architectures for Analog-to-Information Conversion in Channel Equalizers

Aseem Wadhwa*, Upamanyu Madhow* and Naresh Shanbhag†
*Department of ECE, University of California Santa Barbara, CA 93106, USA, {aseem, madhow}@ece.ucsb.edu
†Department of ECE, University of Illinois at Urbana-Champaign, IL 61801, USA, shanbhag@illinois.edu

***Abstract*—As modern communication transceivers scale to multi-Gbps speeds, the power consumption and cost of high-resolution, high-speed analog-to-digital converters (ADCs) become a crucial bottleneck in realizing "mostly digital" receiver architectures that leverage Moore's law. This bottleneck could potentially be alleviated by designing analog front ends for the more specific goal of *analog-to-information conversion* (i.e., preserving the digital information residing in the received signal). As one possible approach towards this goal, we consider a generalization of the standard flash ADC: instead of implementing n bit quantization of a sample by passing it through $2^n - 1$ slicers as in a standard ADC, the slicers are dispersed in time as well as space (i.e., amplitude). Considering BPSK over a dispersive channel, we first show, using ideas similar to those underlying compressive sensing, that randomly dispersing enough one-bit slicers over space and time does provide information sufficient for reliable demodulation over a dispersive channel. We then propose an iterative algorithm for optimizing the design of the sampling times and amplitude thresholds, and provide numerical results showing that the number of slicers can be significantly reduced relative to a conventional flash ADC with comparable bit error rate (BER). These system-level results motivate further investigation, in terms of both circuit and system design, into looking beyond conventional ADC architectures when designing analog front-ends for high-speed communication.**

I. Introduction

Modern communication transceivers are based on "mostly digital" architectures leveraging the low-cost digital computation enabled by Moore's law to implement sophisticated algorithms in digital signal processing (DSP). A crucial component enabling this is the analog-to-digital converter (ADC), which enables translation of the analog received waveform into a digital signal. However, as link speeds and bandwidths scale up to 10s of Gbps, the cost and power consumption of high-resolution ADCs become prohibitive [1]. It is natural to ask, therefore, whether it is possible to relax the ADC requirement to one of *analog-to-information (A/I) conversion* tailored to the communications application, focusing on reliable recovery of the transmitted data, rather than accurate reproduction of the received signal. Specifically, we consider communication over dispersive channels (focusing on the simplest possible setting of BPSK over a wireline channel), and ask whether it is possible to design analog front ends that are more efficient than conventional ADC architectures.

Consider the flash ADC, typically the architecture of choice as we push up the sampling rate. An n-bit flash ADC consists of $2^n - 1$ comparators sampling synchronously. Each sample is compared in parallel against $2^n - 1$ amplitude thresholds, and these $2^n - 1$ bits thus obtained are encoded into a more compact n-bit representation indexing the quantization bin that the sample falls into. However, instead of distributing slicers over space (i.e., amplitude) alone, we ask whether it is fruitful to disperse slicers over time as well as space. Each slicer in such a design produces a single bit, and these bits are fed directly to the digital backend implementing the channel equalizer, thus also eliminating the $2^n - 1$ bit to n bit digital encoder required per sample for the flash ADC. We show through our numerical results that opening up the design space in this fashion, generalizing from a flash ADC to a *space-time slicer architecture,* allows us to reduce the number of slicers required to attain a given link performance.

Summary of results: Our technical contributions can be summarized as follows:
(a) We first show that, if slicers are randomly dispersed in time and space, then, with sufficiently many slicers, we can preserve the L_1 norms of the pairwise differences between waveforms corresponding to a particular bit taking values 0 or 1. While the performance of optimal channel equalization in AWGN is characterized by the L_2 norms of these differences, our result does imply that, if there is no error floor with unquantized observations, then there is no error floor in using generalized space-time slicers. While this provides a sound theoretical underpinning for our proposed architecture, the number of slicers actually required to provide good equalization performance is significantly smaller than estimates from the theory.
(b) We propose an algorithm for optimally selecting the spatio-temporal locations of slicers, and show that this generalized structure leads to a more efficient utilization of slicers relative to a flash ADC, in that fewer slicers are required for a comparable BER. The reduction in the number of slicers (which reduces power consumption of not just the ADC but also the variable gain amplifier that precedes it), and the elimination of the $2^n - 1$ bit to n bit encoder in a traditional ADC (since the bits at the slicer outputs are fed directly to

This work was supported by the Systems on Nanoscale Information fabriCs (SONIC), one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

the equalizer DSP), opens up the possibility for significant power savings. Of course, as discussed in the conclusions, much further work in terms of detailed power estimation is required to evaluate the extent to which the potential power savings can be realized.

Related work: The effect of heavily quantized measurements on communication systems design and performance has received significant attention recently. For non-dispersive channels, the effect of coarse quantization has been studied for the ideal AWGN channel [2], carrier-asynchronous systems [3], [4], and fading channels [5]. Reference [6] discusses channel estimation with coarsely quantized samples. A number of recent papers [7] [8] [9] consider the problem of equalization with low-precision analog front ends, and propose methods for designing ADC quantizer levels. The emphasis in all of these papers remains on designing multiple slicer thresholds for a given sample, rather than dispersing the slicers over time as we allow. The authors in [7] use (bounds on) the *information rate* between the transmitted symbols and the quantized samples as a criterion for designing quantizer levels. They focus on non-contiguous and vector quantization for symbol spaced sampling, and employ the information bottleneck algorithm to approximately maximize the mutual information. We propose a different metric (more sensitive to the uncoded BER) and an iterative algorithm for optimizing our space-time slicer, choosing both sampling times and amplitude thresholds. The advantage of non-uniform thresholds for equalization has been reported in [8][9]. We arrive at similar conclusions, as discussed in Section IV of the paper, but our results are more general with designs involving non-uniform thresholds spread across time. The DSP algorithms used in these references are simple DFE equalizers: 5 feedback taps and a single feedforward tap in [8], which is only applicable to channels with limited precursor intersymbol interference (ISI), and 2-3 feedforward taps with 2 feedback taps in [9]. Further, the numerical evaluations in [8][9] focus on a high SNR regime with very low BERs ($\sim 10^{-12}$). Since we are interested in understanding how far we can push simplification of the analog front end, we consider lower SNR regimes with uncoded BER of the order of $10^{-3} - 10^{-4}$ (which can be still be cleaned up easily with a lightweight channel code), and use optimal BCJR decoding based on the outputs from our space-time slicers.

The proof of our theoretical result on L_1 norm preservation is analogous to that of the Johnson-Lindenstrauss (JL) lemma [10] which provides a theoretical basis for compressed sensing. The result itself appears at first glance to be similar to the *bit-conservation* principle articulated in [11], but the details and implications are completely different. The result in [11] considers signal reconstruction, and can be roughly paraphrased as saying that n 1-bit observations are equivalent to n/k k-bit measurements. In contrast, our result says that n 1-bit measurements are equivalent to n infinite-precision measurements in terms of guaranteeing the feasibility of reliable data recovery in the low-noise regime (albeit with a smaller error exponent).

II. System Description

Let $h(t)$ denote the continuous valued channel impulse response (CIR), the combined response of the channel and the analog preprocessing, just before the analog-to-digital conversion block. Then the continuous valued received waveform is given by

$$x(t) = \sum_{i=0}^{\infty} b_i h(t - iT_s) + w(t) \tag{1}$$

where b_i's are the transmitted bits taking values $\{+1, -1\}$, T_s is the symbol period and $w(t)$ the AWGN process. For a channel of length L symbols, the support of $h(t)$ lies in $[0, LT_s]$. We denote the corresponding discrete-time response, sampled at the symbol rate, by $\mathbf{h} = [h_1, h_2, ..., h_L]^T$. The channels used in this paper for simulations are 20 inch FR4 high-ISI backplane channel $[.0949,\ .2539,\ .1552,\ .0793,\ .0435,\ .0356]$ [9] and a "fatter" channel $[.227,\ .46,\ .688,\ .46,\ .227]$ [12].

Information rate: Samples of the continuous received signal $x(t)$ are used to decode the transmitted bit sequence $\mathbf{b}$. Let $\mathbf{x}_i^j$ denote the vector of samples (these may or may not be quantized) obtained during the interval $[iT_s, jT_s]$. For the special case of symbol spaced sampling, the length of $\mathbf{x}_i^j$ is $j - i + 1$ (the length for general space-time slicers depends on the specific pattern of sampling times used). The information rate between the transmitted bits and the received samples is given by

$$\begin{aligned} I(\mathbf{b}; \mathbf{x}) &= \lim_{n\to\infty} \frac{1}{N} I(\mathbf{b}_1^N; \mathbf{x}_1^N) = \frac{1}{N} \lim_{n\to\infty} \sum_{i=1}^{N} I(b_i; \mathbf{x}_i^N | b_{i-L+1}^{i-1}) \\ &\geq \frac{1}{N} \lim_{n\to\infty} \sum_{i=1}^{N} I(b_i; \mathbf{x}_i^{i+f} | b_{i-L+1}^{i-1}) \end{aligned} \tag{2}$$

Inequality (2), derived in [7], states that the information rate is lower bounded by the average (over the past bits) of the mutual information between the current bit and the measurements over the next few symbols (f), conditioned on the past bits. Numerical results in [7] show that this lower bound becomes a fairly tight approximation for $f = L$ future symbols, where L denotes the length of the CIR.

III. One-bit Measurements with Random Thresholds

In this section we consider the special case of 1-bit measurements. In order to make decisions on the i^{th} transmitted bit b_i (the discussion that follows holds for any bit, hence we drop the index i where convenient), we restrict attention to measurements in the interval $[iT_s, (L + i)T_s]$ affected by this bit. This choice of observation interval is sensible but arbitrary, and our approach applies to other choices as well. The measurements in this interval are also affected by $L - 1$ "past" ISI bits $(b_{i-L+1}, ..., b_{i-1})$ and $L - 1$ "future" ISI bits $(b_{i+1}, ..., b_{i+L-1})$. Denote the noiseless received waveform in this interval by $s(t)$, suppressing the dependence on the desired bit b_i and the ISI bits from the notation. Without loss of generality, we normalize $h(t)$ so that $s(t)$ lies in $[-1, 1]$. The

main result in this section can be paraphrased as follows: for n one-bit measurements, uniformly spaced in time but with thresholds chosen randomly over $[-1, 1]$, if n is large enough, then it is possible (at high SNR) to reliably distinguish between $b_i = +1$ and $b_i = -1$ with one-bit measurements, as long as it is possible to do so with unquantized measurements.

Let $\mathbf{x}_i^{i+L}$ denote the vector of continuous-valued samples obtained by sampling $s(t)$ uniformly, n times, over the observation interval. Fixing the past ISI bits, we partition the noiseless waveforms corresponding to all possible realizations of the future bits into two sets, each of cardinality 2^{L-1}, corresponding to the two possible values of the "tagged bit" b_i: $\mathcal{S}_0 = \{s(t) \text{ s.t. } b_i = -1\}$ and $\mathcal{S}_1 = \{s(t) \text{ s.t. } b_i = +1\}$ for the continuous-time waveforms $s(t)$, and $\mathcal{X}_0$ and $\mathcal{X}_1$ for the corresponding sampled vectors $\mathbf{x}_i^{i+L}$. In order to focus on whether we can reliably demodulate b_i in the face of ISI, we set the noise level to zero . The lower bound (2) on the information rate equals one (implying that the information rate equals one) at this infinite SNR as long as the set of observations generated by the two different values of the desired bit are mutually exclusive; that is, $\mathcal{X}_0 \cap \mathcal{X}_1 = \emptyset$. This always holds for unquantized measurements, as long as a sample is obtained in the first symbol period ($[0, T_s]$) and the corresponding CIR $h(t)$ value is nonzero. This follows from the fact that, since the past bits are fixed, and future ISI bits do not affect the waveform in the interval $[iT_s, (i+1)T_s]$, $b_i = -1$ and $b_i = +1$ result in different samples in the first entry of $\mathbf{x}_i^{i+L}$. This result is also discussed in [13], where the author considers symbol spaced samples and shows that the lower bound (and hence the information rate) goes to one as SNR increases as long as the first element of the discrete time CIR is nonzero. In general, such guarantees cannot be provided for quantized measurements. However, we show that as long as n is *large*, using randomized thresholds for one-bit quantization results in similar behavior as for unquantized observations.

In general (at any SNR), the performance depends on the amount of overlap/separability between the sets $\mathcal{X}_0$ and $\mathcal{X}_1$. For the purpose of our proof, we employ the *normalized* L_1 norm between each pair of elements $\mathbf{x}_0 \in \mathcal{X}_0$, $\mathbf{x}_1 \in \mathcal{X}_1$, defined as follows:

$$\|\mathbf{x}_0 - \mathbf{x}_1\|_1 = \sum_{i=1}^{n} \delta \, |s_0(i\delta) - s_1(i\delta)| \tag{3}$$

where $s_0(t)$ and $s_1(t)$ are the corresponding continuous time waveforms from sets $\mathcal{S}_0$ and $\mathcal{S}_1$ respectively and δ is the sampling interval (for uniform sampling as assumed in this section, $n\delta = LT_s$). The scale factor δ is included for the normalized L_1 norm $\|\mathbf{x}_0 - \mathbf{x}_1\|_1$ to approximate the continuous time L_1 norm $||s_0 - s_1||_1$ (area between the waveforms) as n gets large. We define the minimum normalized L_1 distance between the two sets as follows:

$$d = \min_{\mathbf{x}_0 \in \mathcal{X}_0; \mathbf{x}_1 \in \mathcal{X}_1} \|\mathbf{x}_0 - \mathbf{x}_1\|_1 \tag{4}$$

For unquantized observations, as noted earlier, $\mathcal{X}_0 \cap \mathcal{X}_1 = \emptyset$, and hence $d > 0$.

Let us now consider what happens when we pass the unquantized sampled vector $\mathbf{x}$ through a series of one-bit quantizers, with the ith sample compared to threshold t_i. The vector of thresholds is denoted as $\mathbf{T} = [t_1, t_2, ..., t_n]^T$, and defines a quantization function q as follows:

$$q(\mathbf{x}) = (2\delta)\mathbf{y}$$
$$\mathbf{y}(i) = \begin{cases} 1 & \text{if } \mathbf{x}(i) \geq t_i \\ 0 & \text{if } \mathbf{x}(i) < t_i \end{cases} \quad i = 1, ..., n \tag{5}$$

Thus, $q(\mathbf{x})$ is a scaled version of the binary vector $\mathbf{y}$, where the scaling by δ facilitates comparison of norms before and after quantization.

The following theorem states that, with a sufficient number of samples n, quantized with random thresholds, the quantization function $q(\cdot)$ approximately preserves the L_1 norm of the unquantized differences $\|\mathbf{x}_0 - \mathbf{x}_1\|_1$. This result bears some similarity to the JL-lemma in which random projections preserve the norm for embeddings to lower dimension subspaces [14].

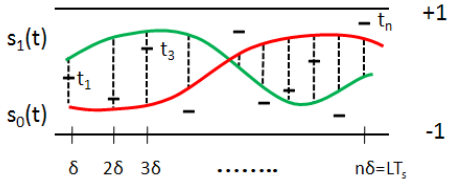

Fig. 1. One-bit Measurements with Varying Thresholds

Theorem. *If each entry of the threshold array* $\mathbf{T}$ *is picked uniformly and independently from* $[-1, 1]$*, then for any constants* $\epsilon, \beta \geq 0$*, with probability at least* $1 - L^{-\beta}$ *, for all* $\mathbf{x}_0 \in \mathcal{X}_0$; $\mathbf{x}_1 \in \mathcal{X}_1$ *we have*

$$(1-\epsilon)\|\mathbf{x}_0 - \mathbf{x}_1\|_1 \leq \|q(\mathbf{x}_0) - q(\mathbf{x}_1)\|_1 \leq (1+\epsilon)\|\mathbf{x}_0 - \mathbf{x}_1\|_1 \tag{6}$$

if n *satisfies*

$$n \geq \frac{4T_s}{d\epsilon^2}\left(2ln2 \cdot L^2 + \beta L lnL\right) \tag{7}$$

where d *is the minimum* L_1 *distance defined in (4).*

Proof:

Consider a particular pair of sampled measurements $\mathbf{x}_0 \in \mathcal{X}_0$; $\mathbf{x}_1 \in \mathcal{X}_1$ (corresponding to $s_0(t) \in \mathcal{S}_0$; $s_1(t) \in \mathcal{S}_1$). Define

$$\mathbf{z} = |q(\mathbf{x}_0) - q(\mathbf{x}_1)| \tag{8}$$

Note that the entries of $\mathbf{z}$ are either 0 or 2δ:

$$\mathbf{z}[i] = \begin{cases} 2\delta \text{ if } t_i \in [\min(s_1(i\delta), s_0(i\delta)), \max(s_1(i\delta), s_0(i\delta))] \\ 0 \text{ otherwise} \end{cases}$$

Thus, an entry is nonzero if the corresponding threshold lies between the two unquantized values. Since t_i is uniformly picked from $[-1, 1]$, $\mathbf{z}[i]$ is a Bernoulli random variable with parameter $p_i = \frac{1}{2}|s_0(i\delta) - s_1(i\delta)|$ and mean $2\delta p_i$. The L_1 norm of $\mathbf{z}$, given by

$$\|\mathbf{z}\|_1 = \sum_{i=1}^{n} \mathbf{z}[i] \tag{9}$$

has expectation

$$E\left(\|\mathbf{z}\|_1\right) = E\left(\sum_{i=1}^{n} \mathbf{z}[i]\right) = \sum_i E(\mathbf{z}[i]) = 2\delta \sum_i \frac{|s_0(i\delta) - s_1(i\delta)|}{2} = \|\mathbf{x}_0 - \mathbf{x}_1\|_1 \tag{10}$$

which follows from Eq. (3). Thus, the quantization function $q(\cdot)$ preserves the norms of the differences in expectation. In order to complete the proof, we derive a Chernoff bound to show that the probability of deviation from the expectation goes to zero for large enough n. To simplify notation, we use the shorthand $u = \|\mathbf{z}\|_1$ and $\mu = \|\mathbf{x}_0 - \mathbf{x}_1\|_1$ in the following.

$$\Pr\left(u > (1+\epsilon)\mu\right) \leq E[e^{au}]e^{-a(1+\epsilon)\mu} \;\; \forall\, a > 0 \tag{11}$$

$$= e^{-a(1+\epsilon)\mu} \prod_{i=1}^{n} E\left(e^{a\mathbf{z}[i]}\right) \tag{12}$$

$$= e^{-a(1+\epsilon)\mu} \prod_{i=1}^{n} \left(p_i e^{2a\delta} + (1-p_i)\right) \tag{13}$$

$$= e^{-a(1+\epsilon)\mu} \prod_{i=1}^{n} \left(1 - p_i(1-e^{2a\delta})\right) \tag{14}$$

$$\leq e^{-a(1+\epsilon)\mu} \prod_{i=1}^{n} e^{-p_i(1-e^{2a\delta})} \tag{15}$$

$$= e^{-a(1+\epsilon)\mu} e^{(e^{2a\delta}-1)\sum p_i} \tag{16}$$

$$= e^{-a(1+\epsilon)\mu} e^{(e^{2a\delta}-1)\frac{\mu}{2\delta}} \tag{17}$$

$$= e^{\mu\left(\frac{1}{2\delta}e^{2a\delta} - \frac{1}{2\delta} - a(1+\epsilon)\right)} \tag{18}$$

Equation (11) follows from the Markov inequality, (12) from (9) and the independence of the thresholds, (15) from the inequality $1 - x < e^{-x}$, and (17) from the expression for p_i and (3). Minimizing the right hand side of Eq. (18) to get the Chernoff bound, we get $a = \frac{1}{2\delta}\ln(1+\epsilon)$. Substituting back, we obtain

$$\Pr\left(u > (1+\epsilon)\mu\right) \leq e^{-\frac{\mu}{2\delta}((1+\epsilon)\ln(1+\epsilon)-\epsilon)} \leq e^{-\frac{\mu n \epsilon^2}{4LT_s}} \tag{19}$$

The desired form in (19) is obtained by using the expansion $\ln(1+\epsilon) = \epsilon - \frac{1}{2}\epsilon^2 + \frac{1}{3}\epsilon^3 - .....$, together with some simplification after substituting $\delta = \frac{LT_s}{n}$. Proceeding along similar lines, we obtain an analogous bound for the probability of deviations below the expectation: $\Pr\left(u < (1-\epsilon)\mu\right) \leq e^{-\frac{\mu n \epsilon^2}{4LT_s}}$. Combining with (19) yields

$$\Pr\left(u < (1-\epsilon)\mu \text{ or } u > (1+\epsilon)\mu\right) \leq 2e^{-\frac{\mu n \epsilon^2}{4LT_s}} \leq 2e^{-\frac{dn\epsilon^2}{4LT_s}} \tag{20}$$

Equation (20) follows from the definition of d (4). We now note that there are 2^{L+1} pairs of distances given the past bits (i.e. $|\mathcal{X}_0| = |\mathcal{X}_1| = 2^L$), and varying the L past bits increases the total number of pairs to be considered to 2^{2L}. In order to obtain the final result we invoke the union bound, which for all $\mathbf{x}_0 \in \mathcal{X}_0$; $\mathbf{x}_1 \in \mathcal{X}_1$ ($u = \|\mathbf{z}\|_1$ and $\mu = \|\mathbf{x}_0 - \mathbf{x}_1\|_1$) and corresponding $\mathbf{z}$ (Eq. 8) gives

$$\Pr\left(\|\mathbf{z}\|_1 \leq (1-\epsilon)\mu \text{ or } \|\mathbf{z}\|_1 \geq (1+\epsilon)\mu\right) \leq 2^{2L} \cdot 2e^{-\frac{dn\epsilon^2}{4LT_s}} \leq L^{-\beta} \tag{21}$$

which can be bounded as tightly as desired (21) by increasing β and ensuring that n meets the condition (7).

■

Remarks: While we have considered uniform sampling for simplicity, this is not required for the theorem to hold. Any non-uniform sampling strategy that provides sufficient density of samples to capture the separation of $s_0(t)$ and $s_1(t)$ in the regions where the waveforms are apart suffices. The continuity of the CIR ensures that the sizes (area) of these separability regions are not insignificant, this property being implicitly captured by the constant d. The independence of the thresholds is a crucial requirement for ensuring that the expected value of norm of the quantized vectors is equal to that of the corresponding unquantized vectors corresponding to the same sampling times.

Simulations: Due to the looseness of the union bound used to prove the theorem, picking n based on the theorem is excessively conservative. We now show via simulations that moderate values of n suffice to provide good equalization performance. Our choice of space-time slicers differs from the set-up of the theorem in two respects:

(1) We pick the thresholds from a Gaussian distribution; this performs far better for moderate values of n than the uniform distribution assumed in the theorem. This is because, while the received signal is scaled to lie in $[-1, 1]$, the density of values near zero is higher (as we vary the possible choices of future ISI bits). For the simulations presented we have used a value of 0.4 for the variance of the gaussian distribution.

(2) Instead of picking n random thresholds over the entire duration of $[0, LT_s]$ corresponding to the span of the CIR, we pick thresholds randomly over a single symbol period T_s. This corresponds to an implementation of slicers operating at the symbol rate with a fixed threshold set for each slicer. This scheme reduces the amount of independence and hence averaging (since the thresholds are now periodic with period equal to the symbol interval), but it is simpler to implement, and provides good BER performance for the channels considered here with 10-20 slicers per symbol.

Figure (2) shows BER curves for two different channels, with optimal Bayesian estimation using the BCJR algorithm [15]. The SNR is defined as $\frac{\|\mathbf{h}\|^2}{\sigma^2}$, where $\mathbf{h}$ is the Nyquist sampled CIR and $\sigma^2$the variance per dimension of AWGN. The BER curves vary slightly for different instances of slicer thresholds, the general behavior remains the same for a fixed number of slicers and we find that ~ 15 slicers are needed to avoid error floors for the first channel, while ~ 20 slicers are required for the "fatter" channel.

IV. Optimizing sampling phases and thresholds

The previous section provides theoretical assurances and simulations when all slicers operate at different sampling instants. At the other extreme, we have a conventional flash ADC, in which all slicers have the same sampling phase. In this section, our goal is to optimize the space-time slicer architecture when we have the freedom to choose both the slicer sampling times and thresholds. We find that such generalized designs perform much better than the two constrained configurations. While we are not addressing details of circuit design, we note that such an architecture could be implemented by controlling the delays of the clock signal driving the slicers (delays of the order of 10s of pico seconds are realizable in current circuit implementations). We constrain the slicers to be sampled at the the symbol rate, which implies the thresholds to be optimized are periodic. Before we present our proposed algorithm for choosing the slicer placements, we discuss a few metrics that could be used to evaluate and compare different slicer designs.

(a) The measure of performance ultimately of interest to us is the *optimal* bit error rate (BER) corresponding to MAP (maximum *a posteriori*) estimation of the transmitted bits. A closed form expression for this BER is not available, hence it must be evaluated numerically by running the BCJR algorithm. The complexity per bit of BCJR grows exponentially with the channel memory and thus it is practical only for short channels (< 10). Moreover, since simulations over long sequences of bits are required to accurately capture the BER up to the order of $10^{-4} - 10^{-5}$, it is prohibitive to run it multiple times over different configurations of slicers. This makes it difficult to directly use the optimal BER as a criterion for optimizing slicer locations.

(b) An alternative metric is the *information rate*, which is used in [7] for designing non-contiguous quantization functions. As with the optimal BER, its numerical evaluation involves running Monte Carlo simulations using the BCJR algorithm [16]. However, this mutual information based measure is not sensitive to BER, and approaches the value of 1 (full capacity) quickly at moderate to low SNRs with relatively high BERs (see simulation plots in [7]). It therefore does not meet our goal of optimizing in the regime of lightly coded systems with low uncoded BERs.

(c) We propose a metric that is simpler to compute than either of the preceding metrics, and is more sensitive to error events than (b). Optimal MAP decoding (required for Monte Carlo estimation of the optimal BER) is difficult because the BCJR decoder must account for a long sequence of received samples for computing posterior probabilities. One way of cutting down on complexity is to cut off the memory required by making hard decisions on past bits, as in the MMSE-DFE. This motivates us to consider the class of sequential equalizers, which make a hard decision on the current bit i based on quantized measurements (denote by $\mathbf{y}$) during $[iT_s, (i+1+f)T_s]$ (i.e. over the current symbol period and f future periods), assuming that reliable hard decisions are available for the past bits (i.e., up to bit $i-1$). While optimal equalization in this class (which we could term the MAP-DFE) is less complex than full BCJR equalization, simulation-based evaluation of the BER as a building block for optimizing thresholds is still too cumbersome. Instead, we define a metric which measures the degree of *statistical overlap* between the observations corresponding to the different values of the current bit. We note that errors in decoding are made when bits 0 and 1 result in similar observations $\mathbf{y}$, and define a metric which computes the probability of this event:

$$\text{average probability of overlap} \propto \sum_{\text{past bits}} \sum_{\mathbf{y}} p(\mathbf{y}|b=0;\ \text{past bits})\ p(\mathbf{y}|b=1;\ \text{past bits}) \quad (22)$$

The metric (22) can be computed with complexity scaling linearly with the number of slicers n_s (in contrast to the exponential complexity in n_s of a single run of BER simulations of the MAP-DFE decoder). For the values of interest ($L, f \sim 4-7$) considered in this paper, the proposed metric can be computed quickly enough to enable repeated evaluations over a space-time grid. It forms the basis for the greedy iterative algorithm discussed next. Due to space limitations, we omit the details of the evaluation of (22) and the complexity analysis.

Iterative Algorithm: The problem of choosing locations for a given number (n_s) of slicers can be broken down as follows: suppose we have already placed k slicers, what is the best spatio-temporal location for the $(k+1)^{\text{th}}$ slicer? To find that we grid the amplitude-time space and for each grid cell compute the metric (c) corresponding to the configuration produced by placing the new slicer at that location. We pick the grid cell that results in the lowest metric and iterate. For our simulations, we have used a fine grid for the amplitude and a coarse one for time (size $T_s/10$, which corresponds to 10 ps for symbol rate of 10 Gbps). At some iterations there may be more than one minimum with the same magnitude, in which case we choose them all to place more than one slicers. Figure (2) shows the final slicer structures obtained for a few different values of n_s and the corresponding BER plots. Note that the configurations obtained are a mixture of the two extreme structures of a flash ADC and 1-bit measurements. Also note the higher concentration of the slicers near the zero threshold (similar to the better performing non-uniform ADC thresholds found in [9][8]). We observe that the generalized structure with 14 slicers results in 3-5dB SNR gains at the BER of 10^{-3} compared to a conventional 4-bit uniform ADC that uses 15 slicers. There are considerable savings in terms of the number of slicers required to achieve a similar performance, we see that 9-10 slicers achieve similar error rates at high SNRs (we still see performance gains at low SNRs). The 1-bit structure (slicers picked randomly as discussed in the previous section) seems to perform marginally better at low SNR (due to a larger number of independent measurements), but is the slowest to attain BERs below 10^{-5}. Thus, despite the suboptimal nature

of the optimization scheme used (our metric is only indirectly related to the BER, and we employ a greedy algorithm), these numerical results on BCJR performance show that carefully designed space-time slicer architectures can lead to more efficient extraction of digital information from the analog signals.

Remarks: Note that we have assumed independence of noise across different time instants for designing the thresholds and performing equalization. This has allowed us to extend the unquantized BCJR algorithm in a straightforward manner for the generalized architecture by treating the outputs of slicers sampling at different times independently. However in an actual system we expect the bandwidth of the frontend filter to be only as wide as the symbol rate which will introduce significant noise correlations across sub-nyquist samples. Simulating the system with correlated noise but using the BCJR algorithm assuming independence results in BER curves for the generalized structure (shown in gray for 14 slicers in Fig. (2)) that are closer to the 4-bit ADC case, for which the BCJR algorithm is still correct. Hence an important topic for future work is to devise a modification of the BCJR that exploits these noise correlations.

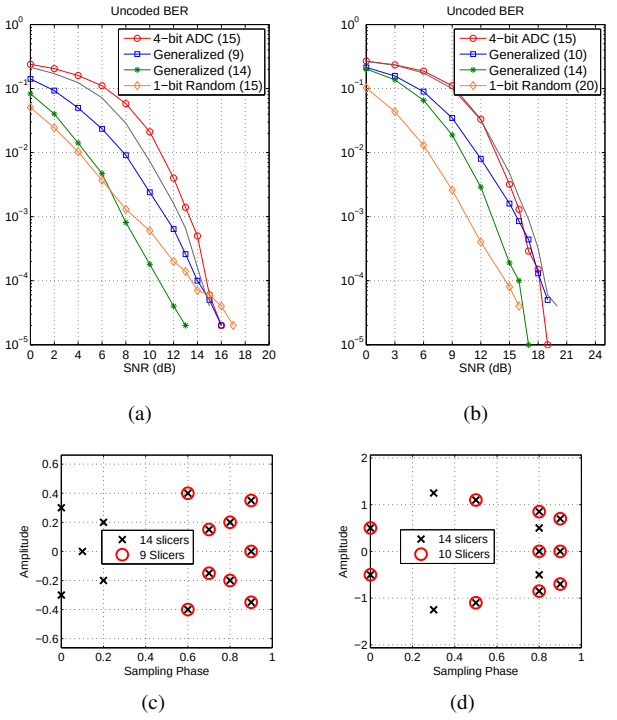

Fig. 2. **(a)** BER plots for channel = $(.0949\ .2539\ .1552\ .0793\ .0435\ .0356)$ **(b)** channel = $(.227\ .46\ .688\ .46\ .227)$ **(c,d)** Generalized slicer locations picked by the iterative algorithm (run at 10 dB) (Values inside the bracket in the legend in subplots (**a, b**) denote the number of slicers (n_s). Curve in gray corresponds to 14 generalized slicers with correlated noise)

V. Conclusions and Future Work

While a dispersive channel increases the dynamic range of a communication signal, we have shown that the information in the transmitted data can be recovered from one-bit slicers dispersed over space and time. For a given BER performance, optimization of this architecture results in a reduction in the number of slicers, and hence potential power savings, relative to a standard flash ADC. Of course, much further design and analysis is required to determine whether the proposed approach indeed results in improved overall performance. For example, the effects of slicer metastability (i.e., uncertainty in digital output when the sample value is close to the threshold) and errors in sampling phases must be accounted for. It is also important to devise low-complexity nonlinear equalizers, since the complexity of the BCJR algorithm used in our results here grows exponentially with channel memory and constellation size.

References

[1] B. Murmann, "ADC performance survey 1997-2013,{Online}." http://www.stanford.edu/∼murmann/adcsurvey.html.

[2] J. Singh, O. Dabeer, and U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *Communications, IEEE Transactions on*, vol. 57, pp. 3629–3639, 2009.

[3] J. Singh and U. Madhow, "On block noncoherent communication with low-precision phase quantization at the receiver," in *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, IEEE, 2009.

[4] A. Wadhwa and U. Madhow, "Blind phase/frequency synchronization with low-precision adc: a bayesian approach," in *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, IEEE, 2013.

[5] G. Middleton and A. Sabharwal, "On the impact of finite receiver resolution in fading channels," in *Allerton Conf. on Communication, Control and Computing*, 2006.

[6] O. Dabeer and U. Madhow, "Channel estimation with low-precision analog-to-digital conversion," in *Communications (ICC), 2010 IEEE International Conference on*, pp. 1–6, IEEE, 2010.

[7] G. Zeitler, A. Singer, and G. Kramer, "Low-precision A/D conversion for maximum information rate in channels with memory," *Communications, IEEE Transactions on*, vol. 60, no. 9, pp. 2511–2521, 2012.

[8] E.-H. Chen, R. Yousry, and C.-K. Yang, "Power optimized adc-based serial link receiver," *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 4, pp. 938–951, 2012.

[9] M. Lu, N. Shanbhag, and A. Singer, "Ber-optimal analog-to-digital converters for communication links," in *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, IEEE, 2010.

[10] S. Dasgupta and A. Gupta, "An elementary proof of the johnson-lindenstrauss lemma," *International Computer Science Institute, Technical Report*, pp. 99–006, 1999.

[11] A. Kumar, P. Ishwar, and K. Ramchandran, "High-resolution distributed sampling of bandlimited fields with low-precision sensors," *Information Theory, IEEE Transactions on*, vol. 57, no. 1, pp. 476–492, 2011.

[12] J. G. Proakis, "Digital communications, 1995."

[13] G. Zeitler, "Low-precision analog-to-digital conversion and mutual information in channels with memory," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, IEEE, 2010.

[14] D. Achlioptas, "Database-friendly random projections," in *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 274–281, ACM, 2001.

[15] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate (corresp.)," *Information Theory, IEEE Transactions on*, vol. 20, no. 2, pp. 284–287, 1974.

[16] D.-M. Arnold, H.-A. Loeliger, P. O. Vontobel, A. Kavcic, and W. Zeng, "Simulation-based computation of information rates for channels with memory," *Information Theory, IEEE Transactions on*, vol. 52, no. 8, pp. 3498–3508, 2006.