This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2007 proceedings.

Communication Limits with Low Precision Analog-to-Digital Conversion at the Receiver

Jaspreet Singh
Dept. of Electrical and Computer Engineering, University of California
Santa Barbara, California 93106, USA
Email: jsingh@ece.ucsb.edu

Onkar Dabeer
School of Technology and Computer Science, Tata Institute for Fundamental Research, Mumbai 400005, India
Email: onkar@tcs.tifr.res.in

Upamanyu Madhow
Dept. of Electrical and Computer Engineering, University of California
Santa Barbara, California 93106, USA
Email: madhow@ece.ucsb.edu

Abstract— **We examine the Shannon limits of communication systems when the precision of the analog-to-digital conversion (ADC) at the receiver is constrained. ADC is costly and power-hungry at high speeds, hence ADC precision is expected to be a limiting factor in the performance of receivers which are heavily based on digital signal processing. In this paper, we consider transmission over an ideal discrete-time real baseband Additive White Gaussian Noise (AWGN) channel, and provide capacity results when the receiver ADC employs a small number of bits, generalizing our prior work on one bit ADC. We show that dithered ADC does not increase capacity, and hence restrict attention to deterministic quantizers. To compute the capacity, we use a dual formulation of the channel capacity problem, which is attractive due to the discrete nature of the output alphabet. The numerical results we obtain strongly support our conjecture that the optimal input distribution is discrete, and has at most one mass point in each quantizer interval. This implies, for example, that it is not possible to support large alphabets such as 64-QAM using 2-bit quantization on the I and Q channels, at least with symbol rate sampling.**

I. Introduction

Modern receiver design is increasingly centered around digital signal processing (DSP), with analog signals being converted to digital format using ADC at the earliest possible opportunity. Operations such as synchronization and demodulation are then performed on the resulting digital signals, greatly increasing the flexibility available to the designer. At the heart of such a design is the implicit assumption that the analog signals can be accurately converted to digital format, by sampling fast enough and quantizing the samples with enough precision (typically 8 bits or more). As communication speeds increase, however, the assumption of adequately accurate ADC may no longer hold: high-speed, high-precision ADC is power-hungry and expensive [1]. In this paper, we ask the question: if we persist with a DSP-centric view of transceiver design, what is the impact of reducing the precision of ADC ? Low-precision ADC impacts every aspect of receiver design, including synchronization, equalization and demodulation. However, our focus here is on characterizing fundamental limits for a discrete-time AWGN channel: the model applies, for example, to ideal Nyquist-sampled linear modulation in AWGN.

[0]This work was supported by the National Science Foundation under grants ANI-0220118 and ECS-0636621, and by the Office of Naval Research under grant N00014-06-1-0066.

We had taken an initial step in this direction when we considered the extreme scenario of 1-bit ADC in [2]. Here we extend it to the general setting of multi-bit ADC and provide results for capacity and optimal input distributions, given a fixed quantizer. For 1-bit ADC, we were able to obtain a closed form expression for capacity. For multi-bit ADC, we do not have analytical expressions for capacity. However, we find that, since the output alphabet is discrete, a dual formulation of the capacity problem provides a systematic methodology for capacity computation. Specifically, we first compute tight upper bounds on the capacity using the dual formulation, and then provide discrete input distributions that very nearly attain these bounds. Moreover, our numerical results support our intuitively plausible conjecture that the cardinality (number of mass points) of these input distributions is bounded by the number of quantization intervals. For example, for a 2-bit quantizer on a real baseband channel, the input alphabet need not have any more than 4 points. Our prior analytical results [2] show that this conjecture holds for the case of 1-bit ADC, for which we proved that BPSK is optimal. Analogous results for the discrete memoryless channel are available in the text by Gallager [3], which shows that the number of input points with nonzero mass need not exceed the cardinality of the output alphabet. Of course, in our scenario, the input alphabet is not a priori discrete, and there is a power constraint, so that the result in [3] does not apply.

In Section II, we show that random dithering cannot improve capacity, allowing us to restrict attention to deterministic quantizers for the rest of the paper. In Section III, we present our method for computing the capacity and near-optimal input distributions, using a dual formulation of capacity that exploits the finite output alphabet. Numerical results are given in Section IV. While we assume a fixed quantizer for our duality-based capacity computations, we attempt to provide numerical insight into quantizer optimization by evaluating capacity for different quantizer choices in Section IV-B.

II. Properties of Optimal Inputs and Quantizers

We consider quantizers for which each bin is an interval of the real line. Therefore a quantizer with M bins is specified by its $(M-1)$ thresholds and the set of quantizers with M

1-4244-0353-7/07/$25.00 ©2007 IEEE

bins can be identified with:

$$\mathcal{Q} := \left\{ \boldsymbol{q} \in \mathbb{R}^{M-1} : -\infty < q_1 < q_2 < \cdots < q_{M-1} < \infty \right\}.$$

The receiver may use dithering, that is, for each symbol duration it may randomly choose some quantizer $\boldsymbol{q} \in \mathcal{Q}$. We assume that the dithering is independent of the channel input and the channel noise. Such random dithering has been found to be beneficial in parameter estimation based on low-precision samples ([4], [5]). Let $\mathcal{P}_{\mathcal{Q}}$ denote the set of probability measures on $\mathcal{Q}$. For a dither distribution $\mu \in \mathcal{P}_{\mathcal{Q}}$, the channel output is given by

$$Y = Q_\mu \left(X + \sigma W \right)$$

where X is the channel input with distribution $F(x)$, W is $\mathcal{N}(0,1)$, and Q_μ denotes a dithered quantizer whose thresholds are chosen as per the probability law μ independently of X and W. Note that we have assumed that the received signal satisfies the Nyquist criteria ([6, pp. 543]), and hence symbol rate sampling of the received waveform results in a discrete memoryless channel. The capacity of this channel under an average power constraint P on the input is

$$C_0 = \max_{\mu \in \mathcal{P}_{\mathcal{Q}}} \max_{F(\cdot):E[X^2] \le P} I(X;Y). \tag{1}$$

Lemma 1: In (1), we can take μ of the form $\delta_{\boldsymbol{q}}$, $\boldsymbol{q} \in \mathcal{Q}$ without any loss of optimality, where $\delta_{\boldsymbol{q}}$ is the unit mass at $\boldsymbol{q}$.

Proof: We note that for a fixed input distribution, $I(X;Y)$ is a convex function of $P(Y|X)$. But

$$P(Y|X) = \int_{\mathcal{Q}} P(Y|X, \boldsymbol{q}) d\mu(\boldsymbol{q})$$

is a linear function of μ. Therefore for a fixed input distribution $I(X;Y)$ is a convex function of μ. To stress the dependence on μ, we write $I(X;Y)$ as $I(\mu)$ now. By Jensen's inequality [7], we have

$$I(\mu) \le \int_{\mathcal{Q}} I\left(\delta_{\boldsymbol{q}}\right) d\mu(\boldsymbol{q}).$$

This implies that for any μ, there exists at least one $\delta_{\boldsymbol{q}}$ such that $I(\mu) \le I\left(\delta_{\boldsymbol{q}}\right)$. Thus without loss of optimality we can only consider deterministic quantizers. ∎

Since the typical quantizers used in practice are symmetric, we further restrict ourselves to symmetric quantizers. Next we show that for symmetric quantizers, without loss of optimality we can only consider symmetric inputs in (1). To be more precise, the input X is symmetric if X and $-X$ have the same distribution, that is, $F(x) = 1 - F(-x)$. For a quantizer $\boldsymbol{q}$, let $\hat{\boldsymbol{q}}$ be the quantizer whose thresholds are the negative of the thresholds of $\boldsymbol{q}$. A symmetric quantizer satisfies $\boldsymbol{q} = \hat{\boldsymbol{q}}$.

Lemma 2: If we only consider symmetric quantizers in (1), then without loss of optimality we can consider only symmetric inputs.

Proof: Suppose we are given $F(x)$. Consider now the following mixture:

$$\tilde{F}(x) = \frac{F(x) + 1 - F(-x)}{2}.$$

This mixture can be achieved by choosing distribution $F(x)$ or $1 - F(-x)$ with probability $1/2$ each. If we use $\tilde{F}(x)$ in place of $F(x)$, we see that $H(Y|X)$ remains unchanged due the symmetric nature of the noise W and the quantizer $\boldsymbol{q}$. However, we note that $H(Y)$ changes. Suppose that when $F(x)$ is used, the PMF of Y is $\boldsymbol{a} = [a_1, ..., a_M]$. Then under $1-F(-x)$ it is $\hat{\boldsymbol{a}} = [a_M, ..., a_1]$. Hence under $\tilde{F}(x)$, the output Y has the mixture PMF $\tilde{\boldsymbol{a}} = (\boldsymbol{a} + \hat{\boldsymbol{a}})/2$. But since the entropy is a concave function,

$$H(Y)\big|_{Y\sim\tilde{\boldsymbol{a}}} \ge \frac{H(Y)\big|_{Y\sim\boldsymbol{a}}}{2} + \frac{H(Y)\big|_{Y\sim\hat{\boldsymbol{a}}}}{2} = H(Y)\big|_{Y\sim\boldsymbol{a}}.$$

It follows that under the symmetric distribution $\tilde{F}(\cdot)$, $I(X;Y)$ is greater than that under $F(\cdot)$. The desired result thus follows. ∎

Even though the use of symmetric quantizers (and hence symmetric inputs) simplifies the problem (1), obtaining an explicit expression for channel capacity is still a formidable task for the multi-bit case. This is due to the complicated expression for the mutual information, which needs to be optimized over a continuous alphabet input. However, the output alphabet is discrete in our problem. In the next section, we exploit this fact by working with an alternate dual formulation of capacity, which involves optimization over the discrete output distribution rather than the input.

III. Methodology for Computing Capacity

In this section, we motivate and outline numerical procedures for obtaining upper bounds on the capacity. We consider a fixed quantizer only; optimization over the quantizer can be carried out by evaluating the mutual information for different quantizers (see Section IV-B). For a fixed quantizer, our problem reduces to finding

$$C = \max_{F(\cdot):E[X^2] \le P} I(X;Y). \tag{2}$$

In principle, we can use the cutting-plane algorithm [8] to find this capacity. However, when the support set of the optimal input distribution is not known, the numerical convergence of this algorithm is extremely slow, and also highly sensitive to initialization. Therefore, in Section III-A we develop a different approach to upper bound the capacity. Further, we conjecture that for our problem the capacity achieving distribution is discrete and in Section III-B we present an approach to find good support sets for the input distribution, which coupled with the cutting-plane algorithm leads to nearly optimal input distributions. In this section we only describe the rationale behind our approach; the goodness of this approach is demonstrated by the numerical results in Section IV.

A. Duality Based Upper Bound on Channel Capacity

The approach we use to derive upper bounds on capacity is motivated by the discrete nature of the channel output, which makes the following dual formulation of capacity more attractive to work with. If we denote the channel transition law (which includes the quantizer operation as well in our case) by $W(y|x)$, and the output distribution by $R(y)$, then channel capacity is given by [9]

$$C = \min_{R(\cdot)} \sup_{x \in \mathcal{X}} D(W(\cdot|x)||R(\cdot)) \tag{3}$$

where $\mathcal{X}$ is the input alphabet, and $D(\cdot||\cdot)$ denotes the relative entropy [7]. Since we know that the output is discrete, working with the above expression seems to be more favorable than using (2), which would require optimization over a continuous alphabet input distribution.

Note that the expression in (3) does not account for the input power constraint P. This is incorporated by introducing a Lagrange parameter, resulting in the following expression for channel capacity ([9], [10])

$$C = \min_{R(\cdot)} \min_{\gamma \geq 0} \sup_{x \in \mathcal{X}} [D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)]. \tag{4}$$

Hence, for every choice of $\gamma \geq 0$ and $R(y)$, we get the following upper bound on channel capacity

$$C \leq U(\gamma, R(y)) = \sup_{x \in \mathcal{X}} [D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)] \tag{5}$$

Optimization over the real-valued x seems difficult to carry out analytically, and it can be carried out numerically (see Section III-C for details of the numerical computation).

We can further minimize the upper bound (5) over γ. For a fixed $R(y)$, we observe that $U(\gamma, R(y))$ is a convex function in γ. This follows from the fact that the upper bound is the pointwise supremum of a set of functions, each of which is convex (in fact affine) in γ ([11, pp. 81]). This observation is vital since we can use it to minimize $U(\gamma, R(y))$ over γ using standard numerical algorithms, giving the following upper bound on capacity for a given output distribution

$$\tilde{U}(R(y)) = \min_{\gamma \geq 0} \sup_{x \in \mathcal{X}} [D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)] \tag{6}$$

Since (6) is a smooth function of $R(y)$, if somehow we can choose the output distribution to be close enough to the one that achieves capacity in (4), then it will provide us a tight upper bound on the capacity. In our case the output distribution is discrete and symmetric (due to symmetric input, symmetric noise, and symmetric quantizer). Hence it can be parametrized by a small number of parameters for low-precision quantizers. For example, for a 2-bit quantizer, the output is characterized by a single parameter. Thus to find good output distributions, one simple way is to search over various values of the output distribution parameters. Alternately, one may employ gradient descent schemes or simulated annealing type methods for finding global minima.

B. Optimal Input Distribution

While the method described above enables us to find an upper bound on the capacity, it does not provide immediate insight into the optimal input distribution, which is required for designing good modulation constellations. In this subsection, we describe a method to find nearly optimal input distributions.

From [12], we know that the following Kuhn-Tucker condition holds for checking the optimality of an input distribution. An input random variable X^* with distribution $F^*(x)$ will achieve the capacity C if and only if there exists a $\gamma \geq 0$ such that

$$\gamma(x^2 - P) + C - D(W(\cdot|x)||R^*(\cdot)) \geq 0 \tag{7}$$

for all $x \in \mathcal{X}$, with equality if x is in the support of X^*, where $R^*(y)$ denotes the output distribution when the input distribution is $F^*(x)$. Using this condition, it is easy to see that if X^* with distribution $F^*(x)$ achieves the capacity (with $\gamma = \gamma^*$ in (7)), then (5) will be satisfied with an equality ($C = U(\gamma^*, R^*(y))$). Hence, $R^*(y)$ will achieve capacity in the dual formulation of (4), with the maximum over x being achieved at those points that belong to the support of X^*.

We now explain how the numerical computation of tight upper bounds using (6) can be leveraged to obtain good approximations for the support set of X^*. Let $\mathcal{X}^*$ be the optimal support set. We note that $g(x, \gamma, R(y)) := D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)$ is a smooth function of all its arguments. Therefore if $R(y)$ is perturbed from its optimal value $R^*(y)$, we expect the solution (x, γ) of (6) to be slightly perturbed from the optimal (x^*, γ^*). Moreover, we expect that if $(\gamma, R(y))$ are slightly perturbed from the optimal values $(\gamma^*, R^*(y))$, then as a function of x, $g(x, \gamma, R(y))$ has local maxima close to $\mathcal{X}^*$. With this motivation, we use the following procedure to find good input distributions.

- For a fixed choice of $(\gamma, R(y))$, find the local maxima of $g(x, \gamma, R(y))$ w.r.t. x.
- Optimize over $(\gamma, R(y))$ as described in Section III-A.
- The set of local maxima (w.r.t. x) corresponding to the optimal solution above is an estimate of $\mathcal{X}^*$. Applying the cutting-plane algorithm, or the Blahut-Arimoto algorithm on this set yields the input distribution optimal for this fixed support set.

Our results in Section IV show that this leads to nearly optimal input distributions.

An alternate method to generate good input distributions is to use the cutting plane algorithm on a very fine grid. Based on the results, we again see that the resulting distributions very nearly achieve the capacity bounds.

C. Details of Numerical Computation

We first note that since the input distribution, noise, quantizer, and (hence) the output distribution are all symmetric, the relative entropy $D(W(\cdot|x)||R(\cdot)) = \sum_y W(y|x) \log \frac{W(y|x)}{R(y)}$ is symmetric in x. The function $D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)$ is thus symmetric, and hence the search for supremum over x in (5) can be confined to just $x \geq 0$.

To perform the maximization over x in (5), we evaluate the function value over a very fine discrete set of x, and pick the maximum. The issue of deciding the range of x over which we evaluate the function is resolved in the following way. For convenience, let us denote $h(x) = D(W(\cdot|x)||R(\cdot))$. It is easy to see that as $x \to \infty$, $h(x)$ saturates (asymptotically reaches a finite value, depending on the output probability distribution

$R(\cdot)$). (The proof is given in the Appendix). However, for all practical purposes, we can assume that this asymptotic value is achieved for $x = k$ where k is sufficiently large. Consequently, the increase in the value of $h(x)$ on going beyond k can be assumed negligible. Combining this observation with the fact that $\gamma(P - x^2)$ is a decreasing function in x, we can conclude that the increase in the value of $h(x) + \gamma(P - x^2)$ (if at all there is an increase) beyond k would also be negligible. Moreover, since the monotonically decreasing function $\gamma(P - x^2)$ dominates the saturating function $h(x)$ for large x, all the local maxima appear inside some bounded set. In our simulations, we search for the local maxima only inside the interval $[0, 10\sqrt{P}]$. We note that because of the power constraint on the input, even if there is a local maxima outside this interval, it will be assigned a negligible probability mass.

The use of a (very fine) discrete grid to perform the maximization over x implies that the computation of the upper bound in (6) can be carried out using the standard Matlab function $fminimax$. The function $fminimax$ uses a a Sequential Quadratic Programming (SQP) method to minimize the worst-case value of a set of multivariable functions, which is precisely what we desire. The fact that it may give local solutions (depending on initialization) is not a concern for us because of the convexity of $U(\gamma, R(y))$ in γ.

IV. Numerical Results

In this section, we present the results when a 2-bit ADC is used at the receiver. In this case, the output distribution is characterized by a single parameter, and therefore we used a simple search technique to find good output distributions. We note that a similar approach may be used for other quantizers, but in general the optimization over the output distribution may be more efficiently done with gradient descent or simulated annealing type methods.

A. Bounds for Fixed Quantizer

A 2-bit symmetric quantizer is characterized by just one parameter. Let us denote the quantizer thresholds by the set $\{-a, 0, a\}$. Since we are using symmetric inputs with a symmetric quantizer, the output is also symmetric, and hence characterized by just one parameter as well. Let the probability distribution on the output be $\{0.5 - \alpha, \alpha, \alpha, 0.5 - \alpha\}$. Then, for a fixed a (fixed quantizer), the dual expression for capacity in (4) becomes

$$C = \min_{0 \le \alpha \le 0.5} \min_{\gamma \ge 0} \sup_x \left[D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)\right] \quad (8)$$

and the corresponding bound in (6) becomes

$$C \le \tilde{U}(\alpha) = \min_{\gamma \ge 0} \sup_x [D(W(\cdot|x)||R(\cdot)) + \gamma(P - x^2)] \quad (9)$$

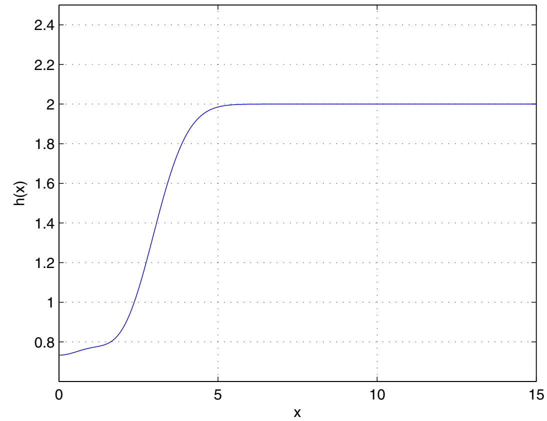

Fig. 1. The relative entropy between channel transition and channel output distributions, plotted against the input.

SNR(dB)	-5.23	-1.15	2.93	7.01	11.10	15.18
Upper Bound	0.156	0.332	0.653	1.097	1.446	1.493
MI_1	0.148	0.330	0.649	1.094	1.428	1.484
MI_2	0.148	0.330	0.650	1.095	1.440	1.484

TABLE I

Achieved mutual information is very near to the upper bound.

where $D(W(\cdot|x)||R(\cdot))$

$$\begin{aligned}
&= \sum_y W(y|x) \log \frac{W(y|x)}{R(y)} \\
&= Q(\frac{a+x}{\sigma}) \log \left[Q(\frac{a+x}{\sigma})\right] + Q(\frac{a-x}{\sigma}) \log \left[Q(\frac{a-x}{\sigma})\right] \\
&\quad + \left[Q(\frac{-x}{\sigma}) - Q(\frac{a-x}{\sigma})\right] \log \left[Q(\frac{-x}{\sigma}) - Q(\frac{a-x}{\sigma})\right] \\
&\quad + \left[Q(\frac{x}{\sigma}) - Q(\frac{a+x}{\sigma})\right] \log \left[Q(\frac{x}{\sigma}) - Q(\frac{a+x}{\sigma})\right] \\
&\quad + \left[Q(\frac{a+x}{\sigma}) + Q(\frac{a-x}{\sigma})\right] \log \left[\frac{\alpha}{0.5-\alpha}\right] + \log \left(\frac{1}{\alpha}\right)
\end{aligned} \quad (10)$$

$Q(x)$ is the complementary Gaussian distribution function:

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp(-t^2/2) dt.$$

With $a = 2$, $\sigma^2 = 1$ and $\alpha = 0.25$, $D(W(\cdot|x)||R(\cdot))$ is plotted in Fig (1). In accordance with the result in the Appendix, we can see that the asymptotic value is $\log(\frac{1}{0.25}) = 2$.

We now present capacity results for the quantizer parameter $a = 2$. We took $\sigma^2 = 1$, and varied the power constraint P to get different SNRs. To compute upper bounds on capacity, the parameter α was varied over a fine grid on $(0, 0.5)$, and for each value of α, the expression in (9) was computed numerically, using the procedure discussed in Section III-C. Table I gives the results we obtained. The numerically computed least upper bound on capacity is depicted first. The next row shows the mutual information (MI_1) achieved by optimizing the input distribution (using cutting plane algorithm) over the

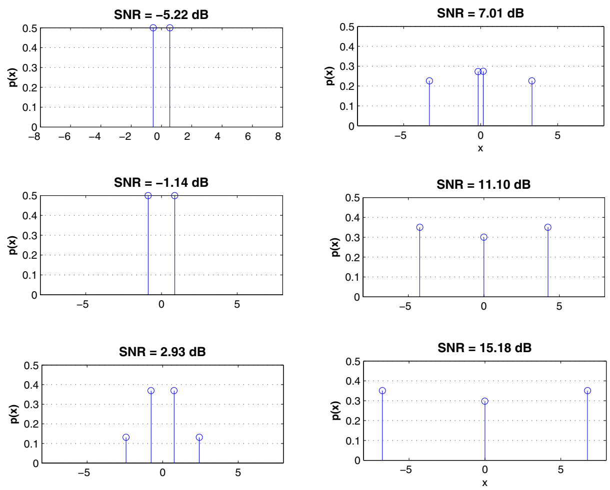

Fig. 2. Probability Mass function of a nearly optimal input at various SNRs

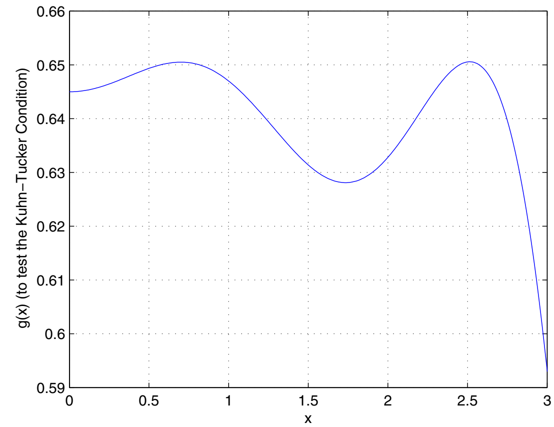

Fig. 3. Near Optimality of the Input Based on the Kuhn-Tucker Condition

support set obtained from the numerical computations to find the upper bounds (as described in section III-B). As is evident from the results, the achieved mutual information is very near to the capacity bound at each SNR.

The numerical results support our conjecture that the capacity achieving distribution is discrete and has at most one mass point in each quantization interval. Fig. 2 illustrates this point, showing the input distributions that achieve the mutual information (MI_1) shown in Table I. We can also see that BPSK achieves the capacity at low SNR as expected [8].

We also employ an alternate direct method to obtain the optimal input distribution (rather than using the approach in section III-B). Specifically, we take a very fine grid in x as the support set and use the cutting plane algorithm to optimize the distribution over it. The results we get using this method are again very close to the capacity bounds, and are shown as MI_2 in the last row of Table I.

It is insightful to look at the Kuhn-Tucker condition (7) for the nearly optimal input distributions obtained above. For instance, let us consider a SNR of 2.93 dB, for which the input random variable X_1 with support set $\mathcal{X}_1 = \{-2.42, -0.76, 0.76, 2.42\}$ and probability mass function $F_1 = \{0.132, 0.368, 0.368, 0.132\}$ achieves a mutual information of 0.6488. If this were an optimal input, then the Kuhn-Tucker condition tells us that there should exist a γ_1 such that the expression $g(x) = D(W(\cdot|x)||R_1(\cdot)) + \gamma_1(P - x^2)$ attains a value equal to the capacity at points in the support set, and a value less than the capacity at all other points, where $R_1(y)$ is the resulting output distribution. However, from the results given in Table I, we know that another input distribution achieves a mutual information of 0.6504 at this SNR, thus implying that (X_1, F_1) is not optimal. Still, if we look at Fig. 3, which shows the plot of $g(x)$ for $x \geq 0$ with γ taken as 0.2115, we see that it has local maxima at $x = \{0.71, 2.52\}$ and the value of $g(x)$ at these points is $\{0.6505, 0.6506\}$. So, even though (X_1, F_1) is not optimal, it still 'almost' satisfies the Kuhn-Tucker condition. This supports our methodology for estimating the optimal support set.

Coming now to the actual capacity values, we observe that even at very high SNR, this channel does not have a capacity of 2 bits, as one would rather expect. It is easy to understand the reasoning behind this. To get a capacity of 2 bits, we need $H(Y) \to 2$ and $H(Y|X) \to 0$. For the conditional uncertainty to tend to zero, the output should be almost deterministic for any input (no cross-overs). In such a situation, to get $H(Y) \to 2$ (all outputs equally likely), the input must have a significant probability mass (ideally one-fourth) in each quantizer interval. However, since $a = 2$ and $\sigma^2 = 1$, a significant input probability mass in the interval $[0, 2]$ would result in a substantial value of $H(Y|X)$, thus implying that it is not possible to ensure that $H(Y) \to 2$ and $H(Y|X) \to 0$ together, and hence the capacity can not approach 2 bits. Based on this observation, we can conclude that for $\sigma^2 = 1$, $a = 2$ is not a good choice for the quantizer if we are in the high SNR regime. The reasoning above suggests that increasing a will improve the chances of $H(Y) \to 2$ and $H(Y|X) \to 0$ simultaneously, and hence the capacity should improve at high SNR by increasing a. This fact is illustrated in the next section, when we see the effects of changing the quantizer on the capacity at different SNRs.

Note that all the results presented here have been obtained for a particular choice of the noise variance, i.e, $\sigma^2 = 1$. However, these results are scale invariant in the sense that if both P and σ^2 are scaled by the same factor K (thus keeping the SNR unchanged), then there is an equivalent quantizer (obtained by scaling the thresholds by $\sqrt{K}$) that gives an identical performance.

B. Optimization Over Quantizer

In this section, we consider the effect of changing the quantizer parameter. This helps us to get some insight about the optimal quantizer for a given SNR. In Fig. 4, we plot the variation of the capacity bounds (which are close to the true capacity) with the quantizer, for different SNR values. Looking at the plots, we observe that as the SNR increases,

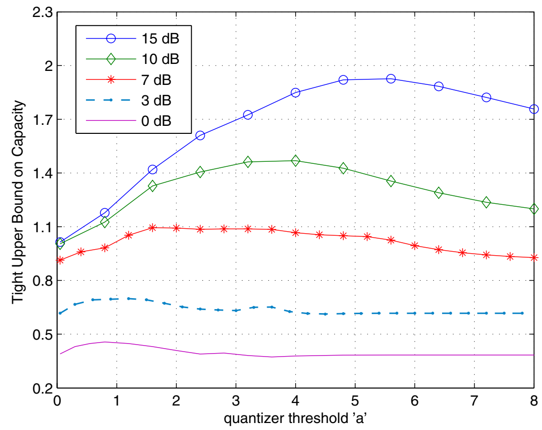

Fig. 4. Optimal quantizer has larger thresholds as the SNR increases (noise variance assumed constant).

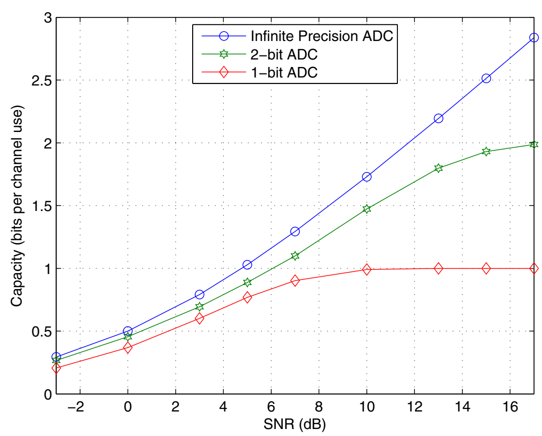

Fig. 5. Capacity with 1-bit, 2-bit and infinite precision ADC.

the optimal quantizer has a larger threshold a. As explained before in section IV, this is understandable since if more signal power is available (for the same noise power), increasing a will make it possible to reduce $H(Y|X)$ and increase $H(Y)$ simultaneously. However, for any finite SNR, the power constraints limit the probability mass that can be put in $[a, \infty)$, which in turn implies that $H(Y)$ is bounded away from 2. Thus increasing a for fixed SNR does not lead to unlimited gains and in general there is an optimum value of a, as depicted in the plots.

Finally, using the results from quantizer optimization, we now plot the capacity with 2-bit quantization. Fig. (5) shows this plot, as well as the capacity plots with 1-bit and infinite precision quantization. We can see that at low to moderate SNRs, the penalty due to low-precision sampling is quite acceptable: for example, at 5 dB SNR, the capacity with 2-bit quantization is 0.89 bits per channel use, as compared to 1.03 bits per channel use with infinite precision. Of course, the number of quantizer bits bounds the output entropy, and hence the capacity.

V. Conclusions

Using the duality based bounds on channel capacity, we have presented a numerical procedure to compute very tight upper bounds on the capacity and generate nearly optimal input distributions for a discrete-time AWGN channel with low precision ADC at the receiver. The results obtained support our conjecture that the optimal input distribution is discrete and has at most one mass point in each quantizer interval. As far as optimization over the quantizer is concerned, we have shown that attention can be restricted to deterministic quantizers without loss of optimality. Developing systematic procedures for quantizer optimization is an important topic of ongoing investigation.

Appendix

Let the quantizer have M bins, with the set of thresholds $\{-\infty = q_0 < q_1 < q_2 < \cdots < q_{M-1} < q_M = \infty\}$. Let the output of the quantizer be denoted by y_i when its input lies in the range $(q_{i-1}, q_i]$, where $1 \leq i \leq M$.

Consider, $h(x) = D(W(\cdot|x)||R(\cdot))$

$$= \sum_i W(y_i|x) \log \frac{W(y_i|x)}{R(y_i)}$$

$$= \sum_i W(y_i|x) \log(W(y_i|x)) - \sum_i W(y_i|x) \log(R(y_i))$$

For any finite noise variance, as $\mathbf{x} \to \infty$, the conditional PMF $W(y_i|x) = \mathsf{P}(Y = y_i|X = x)$ tends to the unit mass at $i = M$. Combining this with the fact that the entropy of a finite alphabet random variable is a continuous function of its probability law, we get

$$\lim_{x \to \infty} h(x) = 0 - \log(R(y_M)) = -\log(R(y_M))$$

References

[1] R. Walden, *Analog-to-Digital Converter Survey and Analysis*, IEEE J. Select. Areas Comm., 17(4):539–550, Apr. 1999.
[2] O. Dabeer, J. Singh and U. Madhow, *On the Limits of Communication Performance with One-Bit Analog-To-Digital Conversion*, In Proc. SPAWC'2006, Cannes, France.
[3] R. G. Gallager, *Information Theory and Reliable Communication*, John Wiley and Sons, New York, 1968.
[4] H. C. Papadopoulos, G. W. Wornell and A. V. Oppenheim, *Sequential signal encoding from noisy measurements using quantizers with dynamic bias control*, IEEE Tran. Info. Theory, 47(3):978–1002, Mar. 2001.
[5] O. Dabeer and A. Karnik, *Signal parameter estimation using 1-bit dithered quantization*, to appear, IEEE Tran. Info. Theory, Dec. 2006.
[6] J. G. Proakis, *Digital Communications*, McGraw Hill, New York, third edition, 1995.
[7] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley Series in Telecommunications, New York, 1991.
[8] J. Huang and S. P. Meyn, *Characterization and Computation of Optimal Distributions for Channel Coding*, IEEE Tran. Info. Theory, 51(7):2336–2351, Jul. 2005.
[9] I. Csiszar and J. Korner, *Information Theory : Coding Theorems For Discrete Memoryless Systems* , Academic Press, 1981.
[10] M. Chiang and S. Boyd, *Geometric Programming Duals of Channel Capacity and Rate Distortion*, IEEE Tran. Info. Theory, 50(2):245–258, Feb. 2004.
[11] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge Univ. Press, 2004.
[12] I. C. Abou-Faycal, M. D. Trott and S. Shamai, *The Capacity of Discrete-Time Memoryless Rayleigh Fading Channels*, IEEE Tran. Info. Theory, 47(4):1290–1301, May 2001.