

Signal Processing for MultiGigabit Communication

P. Sandeep, Jaspreet Singh, and Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106, USA

Abstract—The sophisticated digital signal processing (DSP) at the core of modern communication receivers implicitly assumes the availability of accurate analog-to-digital converters (ADCs). As communication speeds scale up, however, high-rate, high-precision ADCs are either not available, or are too costly or power-hungry. In this paper, we survey our recent results on DSP-centric receiver design with *sloppy* ADCs. For applications requiring limited dynamic range (e.g., small constellations over line-of-sight channels), we consider the performance achievable with ADCs whose precision is drastically smaller (e.g., 1-4 bits) than those in current practice (e.g. 8-12 bits), with a view to characterizing information-theoretic limits as well as developing practical receiver algorithms. For applications requiring larger dynamic range (e.g., large constellations and/or dispersive channels), we consider receiver architectures centered around time-interleaved ADCs, addressing the mismatch between the parallel ADCs for the specific context of their use within a communication receiver, with a view to preventing error floors in receiver performance.

I. INTRODUCTION

Over the last few decades, communication receiver functionality has migrated increasingly from the analog to the digital domain. For example, in typical cellular and wireless local area network (WLAN) transceivers, local oscillators operating in open loop are used to downconvert the received signal, with carrier synchronization, timing synchronization and demodulation all being performed using DSP after analog-to-digital conversion. This approach provides the economies of scale that have propelled mass market wireless and wireline systems. Implicit in this approach is the assumption that the conversion from analog to digital is “accurate enough.” However, this assumption becomes questionable as communication speeds scale up, since accurate analog-to-digital converters (ADCs) at GHz sampling rates or more are either not available, or too costly or power-hungry [1]. On the other hand, DSP-centric receiver architectures remain as attractive as ever as we go to multiGigabit rates, due to the continuing progress of Moore’s law. Thus, we ask the question: how should we redesign communication transceivers if the ADCs we use in the receiver are *sloppy*? In this paper, we briefly survey some recent work resulting from taking the first steps towards answering this question.

While we focus on fundamental questions in communication theory, we note that there are at least three emerging applications of sophisticated signal processing in multiGigabit communication systems. One is ultrawideband (UWB) communication using unlicensed spectrum in the 3-10 GHz band. A second is millimeter wave communication, including unli-

censed use of the 60 GHz oxygen absorption band for short-range links, and semi-unlicensed use of bands above 70 GHz, which avoid oxygen absorption, for long-range links. A third is optical fiber communication, where complex equalization techniques, coherent modulation, and larger constellations, are getting increased attention as electronics gets faster.

We consider two complementary approaches to using sloppy ADC at multiGigabit rates. The first is simply to use drastically lower precision, say using 1-4 bits rather than 8-12 bits as is typical currently. This approach is appropriate when the dynamic range required is limited, e.g., when using small constellations over a near line-of-sight channel. Our goal here is to understand how much performance is lost due to the reduction in precision, to devise algorithms for synchronization and demodulation that are as “digital” as possible, and to understand whether shifting some of the intelligence to the transmitter alleviates the dynamic range requirements at the receiver. The second approach, appropriate when we require larger dynamic range, is to use time-interleaved ADC, where a number of slower sub-ADCs operate in parallel at staggered times to synthesize a faster ADC. The problem here is the mismatch between the sub-ADCs in gain and timing (and more generally, in transfer function). Left uncompensated, such mismatch leads to error floors. However, generic digital mismatch correction for TI-ADCs is not as effective in terms of either cost or performance as mismatch correction in the specific context of a communication receiver. We illustrate this in the context of an OFDM receiver.

Our focus in this paper is to give the big picture. Technical details are available in the publications and submissions (available as preprints from the authors) that we cite.

II. LOW-PRECISION ADC

We first recall and interpret some results obtained last year [2], [3]. Consider a line-of-sight radio link employing linear modulation, with ideal carrier synchronization (no frequency or phase offset between receiver local oscillator and incoming carrier wave) and ideal timing synchronization (Nyquist sampling). Under our assumption of ideal carrier synchronization, we separate out the in-phase (I) and quadrature (Q) components, so that we can restrict attention to a real baseband channel. If the Nyquist-rate samples are now quantized drastically, we obtain the following quantized discrete-time additive white Gaussian noise (AWGN) channel model:

$$Y_k = Q(X_k + N_k) \quad (1)$$

where Y_k are the quantized samples, X_k are the transmitted symbols, N_k is AWGN.

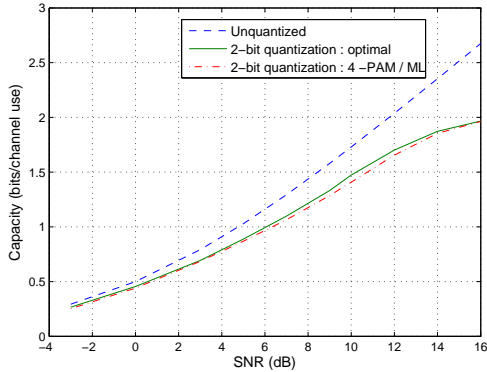


Fig. 1. Capacity with 2-bit quantization: 4-PAM with ML quantization provides close to optimal performance.

The key results from [2], [3], where we investigate the capacity of the quantized AWGN channel (1) under an *average* power constraint on the input, can be summarized as follows: 1) Capacity for the quantized AWGN channel (1) is achievable with a discrete input distribution, with the number of mass points at most $K + 1$, where K is the number of quantization bins. In fact, in all our numerical computations of capacity, at most K mass points suffice.

2) The following intuitively pleasing choice of quantizer and constellation works very well: uniform pulse amplitude modulation (PAM) with quantizer boundaries chosen to coincide with the maximum likelihood decision regions (i.e., chosen midway between the constellation points).

3) The capacity loss relative to unquantized transmission is of the order of 10-15% at moderate signal-to-noise ratio (SNR), which implies that this approach is attractive when power efficiency, rather than bandwidth efficiency, is the primary concern. This is consistent with our goal of supporting higher communication bandwidths.

Given the encouraging nature of these results, the next step is to remove some of the idealizations in the model. We can imagine that timing synchronization should be relatively easy using some form of peak picking, which should be possible if the quantizer provides some amplitude information. For example, 2 bit quantization of the I and Q components should suffice for this purpose. The key bottleneck now becomes carrier synchronization, which we now examine under the assumption of ideal timing synchronization and Nyquist sampling. If the frequency uncertainty is of the order of 100 parts per million (ppm), and the bandwidth is 10% of the carrier frequency, then the data rate is about a 1000 times the maximum frequency offset. For high-precision quantization, the solution is well-established: estimate and correct for the carrier frequency offset in DSP, applying small phase corrections as needed on a per symbol basis in feedback-based tracking mode. Once we drastically reduce ADC precision, however, it is no

longer possible to apply accurate digital phase correction in a feedback loop. Thus, while the rate of phase change due to the carrier offset is typically small relative to the symbol rate (even without any coarse frequency correction), the absolute value of the phase is difficult to track and predict in a heavily quantized system. In order to gain basic insight into this problem, a good model to start with is the block noncoherent model, which in recent years has been studied quite extensively in unquantized settings (e.g., see [4], [5], [6]). In this model, the carrier phase is modeled as constant but unknown over a block of T symbols. For the purpose of capacity computations, we can choose not to use the memory across blocks, and assume for analytical tractability that the carrier phases for different blocks are independent.

The block noncoherent model for a LoS channel can be summarized as follows:

$$Y_i[k] = Q(X_i[k]e^{j\theta_i} + N_i[k]) , \quad k = 1, \dots, T \quad (2)$$

where $\{X_i[k], k = 1, \dots, T\}$ are the symbols in the i th block, $\{N_i[k]\}$ is complex AWGN, θ_i is the unknown phase for the i th block, modeled as uniform over $[0, 2\pi]$, and $Q(\cdot)$ is the quantizer.




Fig. 2. Receiver architecture for 8-sector phase quantization.

Typically, quantizer design is implicitly predicated on automatic gain control (AGC) to ensure that the quantizer levels are set correctly. One-bit ADC, however, does not require AGC. Thus, one approach that is interesting to explore is whether we can use one-bit ADC, along with analog preprocessing, to implement AGC-free quantization. In particular, by taking linear combinations of the I and Q components, and then applying one-bit ADC, we can quantize the phase of the incoming samples, throwing away all amplitude information. Is this enough to handle carrier phase uncertainty, if we restrict attention to phase shift keyed (PSK) constellations? We examine this in the specific context of QPSK, with uniform 8-sector phase quantization. The 8 sectors are obtained by one bit quantization of I, Q, I+Q, and I-Q, as shown in Figure 2.

For QPSK modulation, the receiver architecture in Figure 2 gives rise to a 4-input, 8-output channel with significant rotational symmetries. For example, conditioned on the transmitted symbol, changing the channel phase by $\pi/4$ corresponds to a 1-sector cyclic rotation of the conditional output probabilities. On the other hand, conditioned on the unknown channel phase, changing the transmitted phase by $\pi/2$ corresponds to a 2-sector cyclic rotation of the conditional output probabilities. While these symmetries are intuitively pleasing, they also

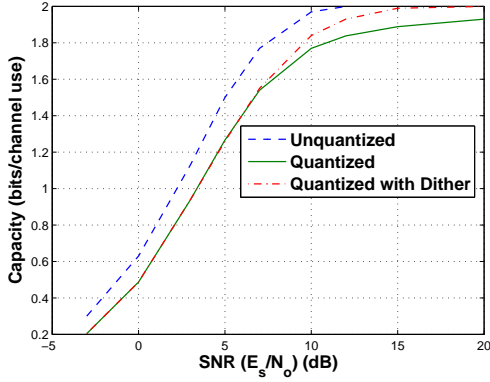


Fig. 3. Capacity for QPSK with 8-sector phase quantization with and without dithering, compared to unquantized performance (block length $T = 4$).

cause problems: it becomes difficult for the noncoherent demodulator to differentiate between phase changes due to modulation and due to the unknown channel phase θ . This causes a fundamental ambiguity: it can be shown that, for certain outputs, the generalized likelihood ratio test (GLRT) demodulator (which jointly estimates the unknown phase and the data) gives two distinct symbol sequences as solutions. This leads to an error floor for uncoded systems, and causes the capacity to rise very slowly to 2 bits/channel use as the SNR increases, which is clearly unacceptable for the high-performance systems that we wish to design. On the other hand, if we dither the transmitted constellation (e.g., by $\pi/8$) on alternate symbols, the ambiguity gets drastically reduced, and both uncoded performance and capacity improve. The relevant capacity plots are shown in Figure 3.

The main conclusion from these preliminary results is that the problem of reliable synchronization and demodulation in the presence of severe quantization at the receiver requires a lot more work, even for a simple LoS channel. The poor performance of the symmetric phase-only quantizer for a symmetric constellation shows that mechanisms for breaking symmetries may be essential. Dithering is one example, but another approach is to include asymmetry in the quantizer design. Finally, amplitude information may be required to enable timing synchronization.

III. TIME-INTERLEAVED ADC

We now discuss the use of time-interleaved ADCs in communication receivers. For communication over a dispersive channel, the unquantized, symbol rate sampled, complex baseband output for a transmitted sequence $\{b[n]\}$ is modeled as

$$y[n] = \sum_k b[k]h[n-k] + w[n]$$

where $\{h[k]\}$ is the discrete-time channel impulse response and $\{w[n]\}$ is noise. For a richly dispersive channel, we would expect $\{y[n]\}$ to begin looking Gaussian by the central limit theorem. In this case, drastic quantization (1-2 bits) is

not expected to give good performance. Accordingly, in this setting, we investigate the use of time-interleaved ADCs of higher precision (e.g., 4-8 bits). The key problem now becomes the mismatch between the component ADCs.

If a signal $x(t)$ is to be sampled by a bank of L ADCs so as to achieve an overall sampling rate of $\frac{1}{T_s}$, then the simplest form of mismatch is in amplitude and timing. Thus, the k th sample of the i th ADC, $i = 1, \dots, L$, is given by

$$x_i[k] = (1 + \epsilon_i)x(kLT_s + \delta_i T_s)$$

where ϵ_i, δ_i are small mismatch errors ($|\epsilon_i|, |\delta_i| \ll 1$). In practice, each ADC may have differences in frequency response which is more complex than can be captured by the simple model above. Two cases of interest are symbol-spaced sampling ($T_s = T$) and fractionally spaced sampling at twice the symbol rate ($T_s = T/2$), with each sub-ADC operating at rate $\frac{1}{2T_s}$.

To develop insight into the effect of timing mismatch, consider a baseband signal $x(t)$ being sampled in parallel by two time-interleaved impulse train samplers to get

$$\begin{aligned} q_1(t) &= x(t + \delta_1) \sum_n \delta(t - 2nT), \\ q_2(t) &= x(t + \delta_2) \sum_n \delta(t - 2nT - T) \end{aligned}$$

In the frequency domain, we have the aliased signals

$$\begin{aligned} Q_1(f) &= \frac{1}{2T} \sum_k X\left(f - \frac{k}{2T}\right) e^{j2\pi\delta_1\left(f - \frac{k}{2T}\right)}, \\ Q_2(f) &= \frac{1}{2T} \sum_k X\left(f - \frac{k}{2T}\right) (-1)^k e^{j2\pi\delta_2\left(f - \frac{k}{2T}\right)} \end{aligned}$$

If there is no timing mismatch, then we have

$$Q_1(f) + Q_2(f) = \frac{1}{T} \sum_k X\left(f - \frac{k}{T}\right)$$

That is, the aliases of $X(f)$ shifted by odd multiples of $\frac{1}{2T}$ cancel out, and we are left only with aliases which are shifted by even multiples of $\frac{1}{2T}$, which corresponds to integer multiples of $\frac{1}{T}$. Thus, when there is no mismatch, the two time-interleaved rate $\frac{1}{2T}$ samplers synthesize a rate $\frac{1}{T}$ sampler, as expected.

When there is mismatch, however, the odd aliased multiples do not quite cancel out. Thus, the desired signal $X(f)$ will be interfered with by aliased versions shifted by multiples of $\frac{1}{2T}$. The dominant terms for this are the aliases $X(f \pm \frac{1}{2T})$.

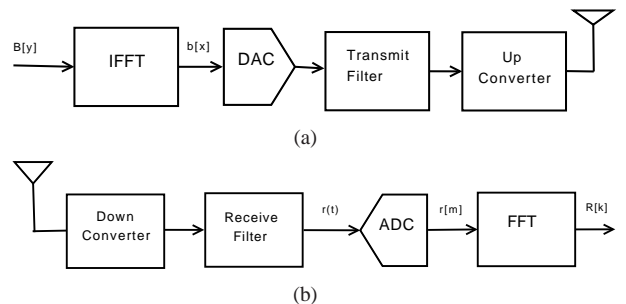


Fig. 4. Standard OFDM transceiver: (a) transmitter (b) receiver

Specifically, consider OFDM with M subcarriers, a standard transceiver for which is depicted in Figure 4. Suppose that

the number of sub-ADCs equals L . Assuming that L divides M , it can be shown that the symbol transmitted on each subcarrier has a nonzero response at exactly L subcarriers. In fact, there are M/L groups of subcarriers, which we may term *interference groups*, such that there is no interference across groups. Within each interference group, we have a classical multiuser detection scenario:

$$\mathbf{y} = \sum_{i=1}^L b_i \mathbf{u}_i + \mathbf{n}$$

where \mathbf{y} is the L -dimensional observation at the receiver corresponding to the interference group, b_i is the symbol transmitted on the i th subcarrier in the group, \mathbf{u}_i is the signal vector corresponding to b_i , and \mathbf{n} is noise. Without mismatch, there is no interference across subcarriers, so that \mathbf{u}_i are orthogonal (each has only one nonzero element). While mismatch destroys this orthogonality, the vectors $\{\mathbf{u}_i\}$ are still near-orthogonal, so that zero-forcing, or decorrelating, reception works extremely well, and produces performance that is close to that with no mismatch. Detailed modeling and results are available in [8].

The complexity of the preceding scheme for mismatch correction scales with the number of parallel sub-ADCs, L , since this is the size of each interference group. If we wish to scale up the communication bandwidth for a fixed sub-ADC speed, however, we would wish to make L large. Fortunately, we have recently discovered a mismatch correction scheme whose complexity does not scale up with L . The key observation here is that the mapping between transmitted samples and received samples, when expressed as a perturbation of the IFFT operation in the OFDM transmitter has only two dominant eigenmodes, so that a zero-forcing approach to mismatch correction can be implemented within the subspace spanned by the two modes. The resulting mismatch model and the corresponding mismatch correction scheme are depicted in Figure 5. Here $B[y]$ denote the frequency domain symbols, $H(y)$ the frequency domain channel, $b[m]$ the time domain transmitted samples, and $r[m]$ the time domain received samples. The dominant modes are denoted by (v_i, u_i) , $i = 1, 2$, and the corresponding singular values are λ_1 and λ_2 . For an oversampled system, it can be shown that the architecture in Figure 5(b) provides a zero-forcing solution, where the time domain coefficients $D[m]$ depend on the specific mismatch values.

While the eigenmode-based approach requires oversampling (unlike the interference group based approach), it has the advantage of scalability, since its complexity is independent of the number of sub-ADCs.

The performance of the architecture in Figure 5 for an OFDM system with 128 subcarriers and 32 parallel sub-ADCs, at an oversampling factor of two, is shown in Figure 6. Thus, each sub-ADC operates 16 times slower than the sampling rate in a standard OFDM system. The gain and timing mismatch levels are both 10%: that is, the gains and timing offsets are chosen uniformly within an interval that varies within $\pm 10\%$ of the nominal values. While Figure 6 shows results

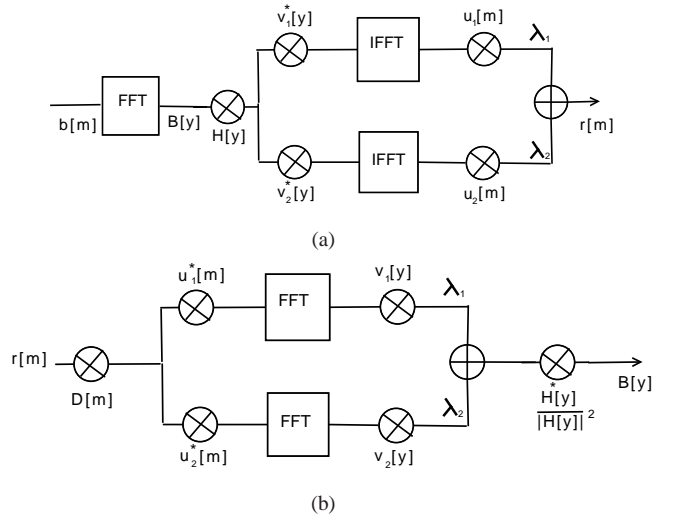


Fig. 5. Scalable architecture for mismatch correction: (a) Approximate model using two eigenmodes; (b) Structure of zero-forcing mismatch correction using the two eigenmode approximation

for a particular realization of mismatch values, the results are typical.

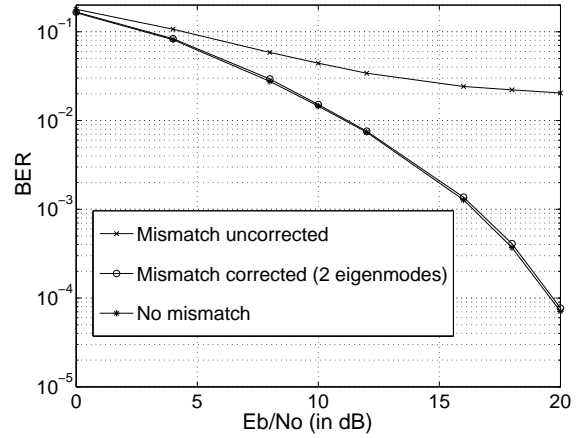


Fig. 6. BER in a 16-QAM, OFDM system (128 sub-carriers) employing a time-interleaved ADC (32 sub-ADCs) with a mismatch level of 10%.

IV. CONCLUSIONS

It is clear from the preliminary results discussed in this paper that we are only at the beginning of our investigation into the impact on sloppy ADC on multiGigabit communication transceiver design. ADC with drastically low precision may have a useful role to play over simple channels, but only if we provide elegant and robust solutions to carrier and timing synchronization, and avoid error floors. Transmit precoding strategies are of particular interest in this context, not only for asymmetric scenarios in which the transmitter is more powerful (e.g, a 60 GHz link from a laptop to a handheld), and also generically, since digital-to-analog conversion at the transmitter scales better in speed than ADC at the receiver. The complementary approach of time interleaving enables

the use of low-power ADC architectures (e.g., pipelined and successive approximation) to build multiGigasample ADCs with enough dynamic range for complicated channels. Our preliminary results show promise for OFDM systems, especially in terms of the scalability of mismatch correction based on the two-eigenmode approximation. However, as we scale the number of sub-ADCs, it is essential to provide effective solutions for channel and mismatch estimation, as well as to devise mismatch correction methods for other settings of interest, including MIMO and singlecarrier systems.

While we investigate of the feasibility of “all-digital” architectures at multiGigabit speeds, we recognize that practical multiGigabit transceiver implementations in the short term may often use analog-centric techniques to circumvent the sloppiness of the ADC. However, there are several reasons to persist with a fundamental investigation of DSP-centric transceiver design with sloppy ADC: first, all-digital architectures are preferable due to economies of scale, hence we must exercise our creativity to the fullest to enable them to the extent possible; second, design with sloppy ADC is a fundamental problem of communication theory which is of intrinsic interest; and third, understanding the limitations of DSP-centric design sheds light on the analog processing that might be needed to support and complement it.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation under grant ECS-0636621 and CCF-0729222. We wish to acknowledge collaborations with our colleagues Dr. Munkyo Seo, Prof. Onkar Dabeer, and Prof. Mark Rodwell.

REFERENCES

- [1] R. Walden, “Analog-to-Digital Converter Survey and Analysis,” vol. 17, no. 4, pp. 539-550, *IEEE J. Select. Areas Comm.*, Apr. 1999.
- [2] J. Singh, O. Dabeer, U. Madhow, “Capacity of the discrete-time AWGN channel under output quantization,” *Proc. 2008 IEEE International Symposium on Information Theory (ISIT 2008)*, Toronto, Canada, July 2008.
- [3] J. Singh, O. Dabeer, U. Madhow, “On the limits of communication with low-precision analog-to-digital conversion at the receiver,” submitted for publication.
- [4] D. Divsalar and M. K. Simon, “Multiple-symbol differential detection of MPSK,” vol. 38, pp. 300-308, *IEEE Trans. on Comm.*, Mar. 1990.
- [5] M. Peleg and S. Shamai, “On the capacity of the blockwise incoherent MPSK channel,” vol. 46, pp. 603-609, *IEEE Trans. Comm.*, May 1998.
- [6] R.-R. Chen, R. Koetter, U. Madhow, D. Agrawal, “Joint noncoherent demodulation and decoding for the block fading channel: a practical framework for approaching Shannon capacity,” vol. 51, no. 10, pp. 1676-1689, *IEEE Transactions on Communications*, October 2003.
- [7] J. Singh and U. Madhow, “On block noncoherent communication with low-precision phase quantization at the receiver,” submitted for publication.
- [8] P. Sandeep and U. Madhow, “Joint channel and mismatch correction for OFDM reception with time-interleaved ADCs: towards mostly digital multiGigabit transceiver architectures,” *Proc. IEEE Globecom 2008*, Nov 30-Dec 4, 2008, New Orleans, Louisiana, USA.
- [9] P. Sandeep and U. Madhow, “Scalable mismatch correction for time-interleaved analog-to-digital converters in OFDM reception,” submitted for publication.