Multi-Gigabit Communication: the ADC Bottleneck[1]

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

***Abstract*—The economies of scale in modern communication systems are enabled by architectures that take advantage of Moore's law to implement most transceiver functionalities in digital signal processing (DSP). The bottleneck in scaling such "mostly digital" architectures to multi-Gigabit rates becomes the analog-to-digital converter (ADC): high-speed, high-precision ADCs are either not available, or are too costly and power-hungry. In this paper, we report on recent results on two approaches towards addressing this bottleneck. The first is to simply use drastically low-precision (1-4 bit) ADCs than current practice. This could be suitable for applications that require limited dynamic range (e.g., line-of-sight communication using small constellations), but there are fundamental and algorithmic questions as to whether all the functions of a communication receiver can be realized with such a significant nonlinearity early in the processing. The second is to use a time-interleaved ADC, where a large number of low-speed, high-precision ADCs are employed in parallel to realize a high-speed, high-precision ADC. This is more generally applicable to applications requiring large dynamic range (e.g., large constellations and/or dispersive channels), but the important question is how to effectively address the mismatch between the component ADCs, which leads to a performance floor if left uncompensated.**

I. Introduction

The ubiquitous adoption of wireless local area networks (WLANs) and cellular networks over the last decade has been propelled by the economies of scale associated with digital receiver architectures that leverage Moore's law for low-power, low-cost implementations. An integral component of such receiver architectures is the Analog-to-Digital Converter (ADC), which converts the received signal into digital format, typically with a precision of 8-12 bits. Operations such as synchronization, equalization, demodulation and decoding are then performed in the digital domain, greatly enhancing the flexibility available to the designer. We would like to scale this "mostly digital" paradigm to multi-Gigabit speeds, in order to enable mass market multi-Gigabit WLANs and wireless personal area networks (WPANs) based on large amounts of unlicensed bandwidth available for Ultra-wideband (UWB) communication (3.1-10.6 GHz in the US) and millimeter wave communication (57-64 GHz in the US). The bottleneck in such scaling becomes the ADC: high-speed, high-precision ADC is either unavailable, or is too costly and power-hungry [1]. In this paper, we discuss two approaches for circumventing this bottleneck by relaxing the requirements on the ADC.

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

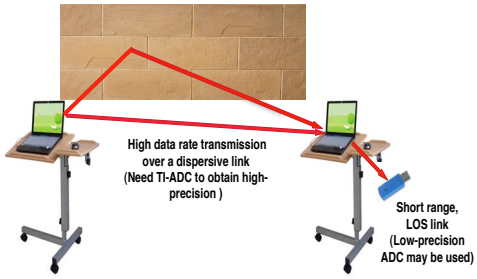

Fig. 1. An in-room WPAN scenario : Low-precision ADC may be suitable for short range LoS communication between a handheld and a laptop. For more demanding applications (laptop-laptop), TI-ADCs could be used to provide larger dynamic range.

The first approach corresponds to simply employing ADCs with much lower precision (e.g., 1-4 bits) than current practice. This approach may be well matched to applications requiring smaller dynamic range, such as line-of-sight (LoS) links and small constellations. An example is the short-range LoS link from laptop to handheld shown in Figure 1, where we might wish to reduce the cost and power consumption of the handheld receiver, possibly at the cost of increasing the transmit power and complexity at the laptop. However, the drastic reduction of ADC precision raises fundamental questions, at both information-theoretic and algorithmic levels, regarding whether it is even possible to engineer a link with such a significant nonlinearity so early in the processing. The second approach corresponds to employing time-interleaved (TI) ADCs, in which a number of low-speed, high-precision ADCs operate in parallel to synthesize a high-speed, high-precision ADC. This is a more generally applicable approach, and is well matched to settings where a larger dynamic range is required. An example is the longer range laptop-to-laptop link shown in Figure 1, where we may wish to use large constellations over a dispersive, possibly multiple-input multiple-output (MIMO), channel. However, mismatch between the component "sub-ADCs," if left uncompensated, ultimately limits communication performance, and the key question is whether the mismatch can be estimated and compensated for effectively as we scale up the number of sub-ADCs.

We begin by providing a brief overview of the state-of-the-art for ADC technology in Section II. In Section III, we discuss link design with low-precision ADC. We begin with information-theoretic limits for an idealized AWGN channel with perfect carrier and timing synchronization, which indicate that low-precision ADC may be an interesting

design option for high-bandwidth communication operating at low to moderate signal-to-noise ratio (SNR). We then take initial steps towards removing the idealization in our model by considering a system without *a priori* carrier synchronism. Link design with TI-ADCs is discussed in Section IV in the context of an Orthogonal Frequency Division Multiplexed (OFDM) system. We show that mismatch between the sub-ADCs leads to interference between subcarriers, and discuss two methods for alleviating this interference that provide better complexity-performance tradeoffs than mismatch compensation schemes devised for generic applications of TI-ADCs. Section V contains our conclusions.

Our intent in this paper is to provide a survey of our prior and new results, along with a discussion of their implications. Further technical details may be found in the cited references, which include preprints available from the authors upon request.

II. ADC Technology

Walden's survey in 1999 [1] provided many useful insights into the technological advances that took place in the field of analog-to-digital conversion over the decade spanning the 90s. Among the notable conclusions from his survey, one was the observation that for a given sampling frequency, it takes roughly 5 years to achieve an improvement of 1 bit in precision. However, as the sampling rates are increased, a major limiting factor in achieving high-precision ADC is aperture jitter, which is the error due to the sample-to-sample variation in the instant of conversion. Specifically, Walden observed that for large sampling rates (above 2 MS/s), the precision falls off by 1 bit for every doubling of the sampling rate. Lundberg's survey [2] in 2002 did not show any significant improvements. The situation today shows the same trends, with the available precision being severely constrained at high sampling rates. In short, ADC technology is not scaling up at a rate commensurate with that demanded by emerging multiGigabit systems, thus making it necessary to explore system design with relaxed ADC specifications.

III. Transceiver Design with Low-Precision ADC

A systems engineer designing DSP-centric transceiver architectures usually takes the following two-step approach:

- Design the system assuming that the ADC at the receiver front-end is perfect (that is, it has infinite precision).
- Perform simulation tests to obtain the minor algorithmic tweaks needed to accommodate finite precision (typically 8-12 bits).

This paradigm for design and implementation is predicated on the assumption that the performance with 8-12 bit quantization essentially replicates that with infinite precision. However, for drastically quantized systems (1-4 bits), the paradigm clearly breaks down, mandating a complete rethinking of system design. To this end, our first step is to characterize the fundamental limits on communication imposed by the use of low-precision ADC. Specifically, how much do we lose in terms of channel capacity by using low-precision ADC? How does the precision of the ADC impact the choice of modulation we use?

We begin by recalling some results we obtained last year while investigating these issues for the classical bandlimited Additive White Gaussian Noise (AWGN) channel [3], [4]. Apart from its fundamental significance, the AWGN channel model also forms an excellent approximation for the near line-of-sight, 60 GHz "point-and-shoot" links, where the use of directional antennas (possibly fixed beam) at each end can cut down drastically on multipath. In doing our capacity analysis, we assume ideal carrier synchronization (no frequency or phase offset between the incoming carrier wave and the local oscillator at the receiver), and also ideal timing synchronization between the transmitter and receiver clocks. In Section III.B, we focus on possible implicit and explicit approaches to attain synchronization.

A. Shannon-Theory for Ideal AWGN Channel

Under our assumption of ideal carrier synchronization, we separate out the the in-phase (I) and quadrature (Q) components, and restrict attention to a real baseband AWGN channel only. We consider linear modulation, with the symbol rate Nyquist samples quantized using a low-precision ADC. This results in the following discrete-time memoryless quantized AWGN channel:

$$Y = Q(X + N) \ , \tag{1}$$

where Y is the quantized output, X is the transmitted symbol, and N is AWGN.

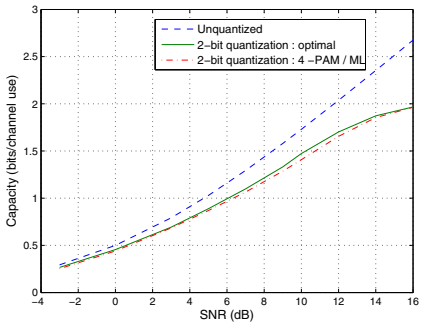

Fig. 2. Capacity with 2-bit quantization: 4-PAM with ML quantization provides close to optimal performance.

In [3], [4], we performed a Shannon-theoretic capacity analysis for this channel, which included an optimization over the choice of the input modulation, as well as over the choice of the quantization thresholds. The results are encouraging: the capacity loss relative to unquantized transmission is of the order of 10-15% at moderate signal-to-noise ratio (SNR), and uniform pulse amplitude modulation (PAM) with quantizer boundaries chosen to be the maximum likelihood (ML) decision regions (i.e., chosen midway between the constellation points) is close to optimal. Example results for 2-bit quantization are shown in Figure 2.

We can now turn our attention to trying to achieve these information-theoretic limits for multi-Gigabit systems. Given

that turbo-like codes approaching Shannon capacity are now available for a wide variety of rates and channels, the key issue is not choice of code, but rather whether we can devise receiver algorithms for timing and carrier synchronization using low-precision ADC. We limit our discussion here to carrier synchronization, which we believe is the more difficult problem.

B. Handling Carrier Asynchronism

Consider linear modulation over a LoS link with ideal timing synchronization. The receiver downconverts the received passband signal with a fixed local oscillator which has an offset Δf relative to the frequency of the incoming carrier. Prior to quantization, the Nyquist-rate complex baseband samples are given by

$$y[n] = b[n]e^{j(n\phi+\theta_0)} + w[n]$$

with $\{b[n]\}$ denoting the transmitted symbols, $\{w[n]\}$ complex AWGN, $\phi = 2\pi\Delta f T_s$ is the phase change per symbol (here $\frac{1}{T_s}$ is the symbol rate), and θ_0 is the initial phase offset. Both ϕ and θ_0 are *a priori* unknown. Now, if we quantize $y[n]$ to obtain $z[n] = Q(y[n])$, how well can we recover the symbols $\{b[n]\}$?

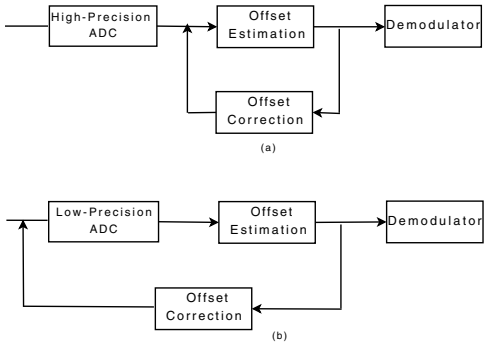

Fig. 3. Feedback-based correction for frequency offset. (a) For high-precision ADC, the correction is done in the digital domain. (b) For low-precision ADC, it may be possible to perform analog offset correction based on digital feedback.

Ignoring quantization effects, the solution is conceptually straightforward: estimate ϕ and θ_0, and estimate the transmitted symbols based on the derotated observations $\{y[n]e^{-j(n\phi+\theta_0)}$. Indeed, this is essentially what is done in conventional DSP-centric architectures based on high-precision ADC, as shown in Figure 3(a), with open loop frequency/phase estimation or feedback-based phase tracking, or a mixture thereof. However, this approach runs into trouble when we drastically lower ADC precision. For example, for a frequency uncertainty of 100 parts per million (ppm) and a bandwidth equal to 10% of the carrier frequency f_c, we have $\Delta f \in [-10^{-4}f_c, 10^{-4}f_c]$, and $\frac{1}{T_s} \approx 0.1f_c$, so that the phase change ϕ per symbol lies in $[-0.006, 0.006]$. Can we estimate and correct for such small rates of phase change when the receiver employs low-precision ADC, especially in a packetized system in which ϕ may change across transceivers? This may require analog offset compensation using feedback generated by post-ADC DSP, as shown in Figure 3(b). Such explicit correction is attractive because it allows us to invoke our results from the last subsection, and to use the near-optimal and *elegant* strategy of uniform PAM with ML quantizers on both the I and Q channels. For example, if we constrain ourselves to 3-bit precision ADCs on the I and Q channels, then a 64-QAM constellation with the ADC thresholds set to implement ML detection will result in only about 10-15 % loss compared to a system using infinite precision quantizers.

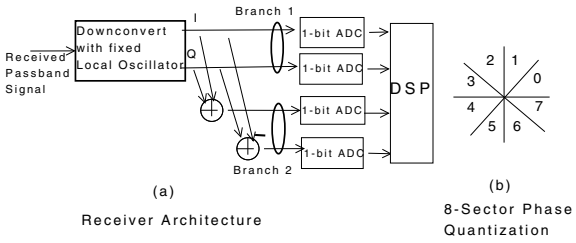

Fig. 4. 8-sector phase quantization using two branches: 1-bit quantization on I and Q on the first branch, and 1-bit quantization on a $\pi/4$ phase rotated version obtained by taking the linear combinations I+Q and I-Q.

We leave devising explicit phase correction architectures as an important open issue, and limit further discussion to an alternative approach: instead of explicit derotation, we can employ *noncoherent* or *differentially coherent* demodulators based on modeling the phase as unknown but constant over a number of symbols. Classical differential modulation and demodulation employs a block of two symbols, but its performance can be improved significantly (especially for larger constellations) by considering larger blocks, at least for unquantized observations (e.g., see [7], [8], [5]). Does this *block noncoherent* approach, which amounts to implicit offset estimation and correction, yield good performance for heavily quantized observations?

1) Block Noncoherent Communication: The quantized block noncoherent channel is modeled as follows:

$$Y_i[n] = \mathsf{Q}\left(X_i[n]e^{j\theta_i} + N_i[n]\right) \ , \quad n = 1, ..., L \qquad (2)$$

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

AGC-Free ADCs: In order to further simplify receiver design, we examine whether acceptable performance can be obtained by removing the need for automatic gain control (AGC). Specifically, by linearly combining the I and Q components in the analog domain, followed by 1-bit ADC (which requires no AGC), we can implement phase-only quantization. An example is the 8-sector phase quantizer shown in Figure 4, where we employ four 1-bit ADCs: two for I and Q, and two for the linear combinations I+Q and I-Q corresponding to a $\pi/4$ phase rotation. Intuitively, we expect phase-only quantization to work well with phase shift keyed (PSK) constellations, but there are potential pitfalls, as we illustrate next for QPSK and 8-sector phase quantization.

Information-theoretic results for phase-only quantization are reported in [6]. Here we report uncoded performance of

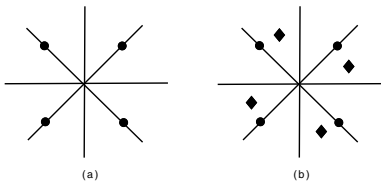

Fig. 5. (a) Standard PSK : the same constellation (the one shown) is used for both symbols in the block. (b) Dithered-PSK : the constellations used for the two symbols are not identical, but the second one is a dithered version of the first one.

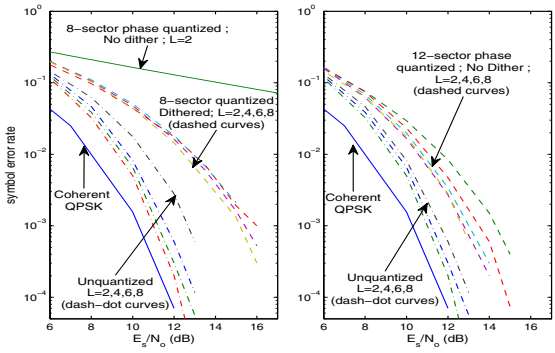

Fig. 6. Symbol error rate performance for QPSK with 8-sector phase quantization (left figure) and 12-sector phase quantization (right figure), for block lengths varying from 2 to 8. Also shown for comparison are the curves for coherent QPSK, and noncoherent unquantized QPSK.

maximum likelihood (ML) block demodulation. For QPSK with 8-phase quantization as in Figure 4, it turns out that there are certain symmetries that have dire consequences. Specifically, these make it impossible to clearly distinguish between the effect of the unknown phase offset and phase modulation. As a result, there are always two ML solutions for certain outputs: as SNR increases, the probability of these outputs decreases, but rather slowly, resulting in poor performance.

One possible way to break the undesirable symmetries is to use dithering. This can be done either at the receiver by using analog pre-multipliers to dither the quantization boundaries, or equivalently at the transmitter by dithering the QPSK constellation from one symbol to the next. Figure 5 depicts a transmit dither scheme, in which the QPSK constellation on the second symbol (shown by diamond shape) is dithered by $-\pi/8$ from the constellation on the first symbol. A careful analysis shows that using such a dithered-QPSK scheme with 8-sector quantization eliminates the ambiguity in ML estimation.

Figure 6 shows the performance curves for QPSK with 8-sector quantization (left plot), and 12-sector quantization (right plot). Also shown for reference is the performance with unquantized block demodulation. For 8-sector quantization, dithering results in substantial performance improvement, but increasing the block length does not provide significant gains. At symbol error rate 10^{-3}, 8-sector quantization with $L = 8$ results in a loss of about 4 dB compared to unquantized observations. On the other hand, with 12-sector quantization, the loss is reduced to about 2 dB, and dithering is not required.

These preliminary results indicate that, while much more invention and evaluation is needed, low-precision ADC may be a feasible option for mostly digital architectures over simple LoS links.

IV. Time-interleaved Analog-to-Digital Converters (TI-ADC)

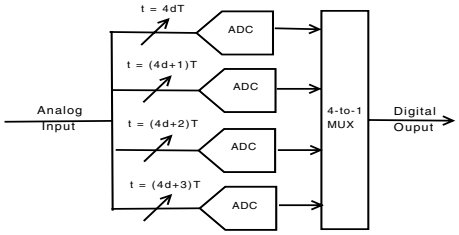

Fig. 7. Ideal TI-ADC with 4 sub-ADCs (d =integer)

Fig. 7. shows a TI-ADC with 4 sub-ADCs interleaved in time. In general, we consider L sub-ADCs, with the goal of scaling L up significantly: slowing the sub-ADCs down allows the use of more power-efficient pipelined or successive approximation architectures, as compared to the flash architectures needed for a direct implementation of a high-speed ADC. Now, effective compensation of the mismatch between the sub-ADCs becomes key to avoiding performance floors in the communication system. Generic mismatch compensation schemes employ long equalizers and are computationally intensive [14], [15], [16]. Our goal is to investigate whether more effective mismatch compensation schemes can be devised by exploiting the specific structure of communication signals. We illustrate this for the example of an OFDM system, reviewing results reported in [12] and previewing results from [13].

A. System model

To a first order approximation, the mismatch between the parallel branches is assumed to be gain, timing and voltage-offset mismatches [9]. Denoting the analog input to the ADC as $r(t)$, the digital output of the TI-ADC with L sub-ADCs is given as,

$$r[m] = g_{m \bmod L} r(mT_o + \delta_{m \bmod L}) + \mu_{m \bmod L} + q[m] \qquad (3)$$

where the *nominal sampling rate* without mismatch is denoted as T_o^{-1}. The m^{th} output sample $r[m]$ is obtained from the sub-ADC with index $m \bmod L$, where $\bmod L$ indicates the remainder after division by L. The gain, timing and voltage offset mismatches of the sub-ADC, with index m, are denoted as (g_m, δ_m, μ_m). Assuming high-precision output, we neglect the quantization noise $q[m]$. We also neglect voltage-offset mismatches in our analysis owing to their simple additive and signal independent structure. As the communication system, we consider the DSP-centric OFDM transceiver [11]. We now discuss two models for mismatch-induced interference which lead to two different compensation structures.

B. Multi-user model

As shown in [12], TI-ADC mismatches lead to interference among the OFDM subcarriers. Figure 8 shows the relative

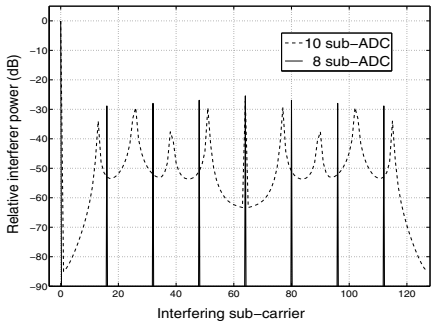

Fig. 8. Effect of TI-ADC mismatch on the OFDM system: Relative interference seen at the frequency bin corresponding to the first subcarrier in 128-subcarrier OFDM system. The values in the legend indicate the number of sub-ADCs interleaved in the TI-ADC. Gain and timing mismatches are uniform in the range $\pm 10\%$ from their ideal values. [12]

interference power at the first sub-carrier due to other subcarriers. In general, all subcarriers interfere with each other, but when L, the number of sub-ADCs, divides M, the number of subcarriers, each subcarrier incurs interference from, and interferes with, precisely $L-1$ other subcarriers. We therefore assume that L divides M, in which case the subcarriers fall into $K = \frac{M}{L}$ mutually disjoint interference groups of size L, so that mismatch compensation can be applied separately to each group. For the j^{th} interference group, $j = 0, 1, ..., \frac{M}{L} - 1$, we obtain a multiuser model as follows:

$$\underline{\mathbf{R}}_j = \sum_{i=0}^{L-1} B_{j+iK} \underline{\mathbf{U}}_{j+iK} + \underline{\mathbf{N}}_j \tag{4}$$

where we denote the received OFDM frame after FFT by $\underline{\mathbf{R}}_j$, the additive noise by $\underline{\mathbf{N}}_j$, and the signal vector corresponding to symbol B_i carried on subcarrier i by $\underline{\mathbf{U}}_i$. We can now apply standard multiuser detection techniques to demodulate the symbols within an interference group. In particular, the signal vectors $\{\underline{\mathbf{U}}_i\}$ are near-orthogonal for moderate mismatch levels, so that the zero-forcing, or de-correlating, detector is found to work very well, incurring minimal noise enhancement. The interference group approach works well for small to moderate values of L, but the equalizer size scales linearly with L. Thus, when a large number of sub-ADCs are to be employed, we must turn to alternative approaches.

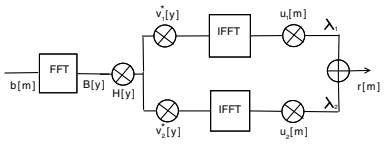

Fig. 9. Approximate model for a TI-ADC, with gain and timing mismatches, in an OFDM system

C. Time-frequency multiplier model

In our attempts to obtain mismatch compensation schemes scalable for large L, we have arrived at an interesting representation for a TI-ADC with gain and timing mismatches [13]. The following approximate relationship holds between the vector of frequency domain symbols $\mathbf{B}$ and the vector of time domain samples $\mathbf{r}$ [13]:

$$\mathbf{r} = A(\overline{\mathbf{H}} \cdot * \overline{\mathbf{B}}) + \overline{\mathbf{n}} \tag{5}$$

where the vector of channel gains is denoted by $\overline{\mathbf{H}}$, the element-wise multiplication by $\cdot *$ and the vector of noise samples by $\overline{\mathbf{n}}$. The elements of the MxM matrix A are given by $A(m, y) = g_m e^{\frac{j2\pi y(m+\delta_m)}{M}}$. For no mismatch, A is the standard IFFT matrix F^*. Consider the matrix $\Delta = A \cdot / F^*$, where $\cdot /$ denotes the element-wise division. Even for large mismatch (deviations as large as 25% in gain and timing from nominal values), we find that Δ is well approximated by the first two terms in its singular value decomposition: there is one significant singular value, with other singular values decaying exponentially to zero.

Using the first two singular values, we can write $\Delta \approx \lambda_1 \mathbf{u_1} \mathbf{v_1^*} + \lambda_2 \mathbf{u_2} \mathbf{v_2^*}$. This implies

$$A = \lambda_1 D(\mathbf{u_1}) F^* D(\mathbf{v_1^*}) + \lambda_2 D(\mathbf{u_2}) F^* D(\mathbf{v_2^*}) \tag{6}$$

The operation of A can be graphically represented as in Fig. 9. It can be seen that the action of a TI-ADC with gain and timing mismatches can be summarized as parallel additive blocks with each block representing multiplication in both time and frequency domains.

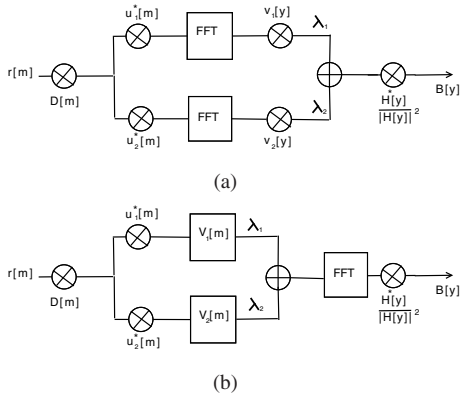

Fig. 10. {(a) Double FFT-based implementation (b) Single FFT based implementation } for time-frequency multiplier based equalizer to compensate TI-ADC mismatch in an OFDM system. Here, $D[m]$ is chosen appropriately so as to cancel all the mismatch-induced interference [13]. In (b), cyclic convolution is performed with the filters indicated by $V_i[m]$ which are obtained by taking IFFT of the vectors $v_i[y]$.

The preceding model motivates the equalizer structure in Fig. 10 (a); see [13] for details. An important feature is that the complexity of this scheme is independent of mismatch levels, the required digital precision and the number of sub-ADCs. Complexity can be further reduced for the alternate representation in Fig. 10(b), which requires only one FFT. It can be shown that the effect of mismatch can be completely eliminated as long as the overall sampling rate is at least twice the Nyquist rate. For an oversampling factor of two, our numerical evaluations show that the filters in Fig. 10(b) have very few significant taps.

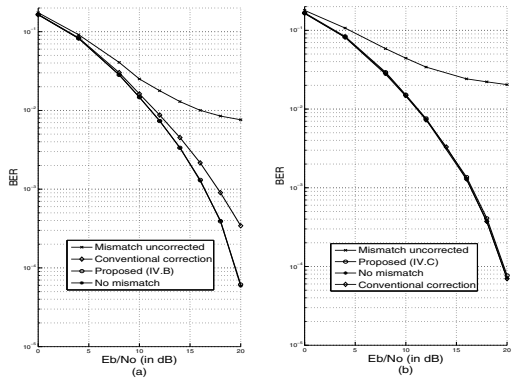

Fig. 11. BER in a 16-QAM, OFDM system (128 subcarriers) employing a TI-ADC with 8 sub-ADCs (a) and 32 sub-ADCs (b), with a mismatch level of 10%.

D. Numerical Results

We now apply the proposed OFDM-specific mismatch compensation algorithms to an OFDM system employing 16-QAM over a dispersive channel. Two values of interleaving factor, moderate (L=8) and large ($L = 32$), are considered. In both cases, 10% gain and timing mismatches are assumed between the interleaved ADCs. See [12], [13] for further details. We compare the performance of the proposed schemes with a conventional mismatch compensation approach of employing a zero-forcing (time varying) equalizer on the time domain output of the TI-ADC [10].

Figure 11(a) depicts the BER performance before and after mismatch correction for 8 sub-ADCs. It can be seen that the method in Section IV.B is effective in eliminating mismatch-induced error floors, and outperforms conventional mismatch correction. The number of taps in the conventional equalizer was chosen to be L(=8), so as to match the complexity of the proposed scheme. Figure 11(b) depicts the performance for 32 sub-ADCs with an aggregate sampling rate equal to twice the Nyquist rate. Using the time-frequency multiplier scheme in Fig. 10(b) (Section IV.C), with the number of taps in $V_1[m]$ and $V_2[m]$ as 1 and 2, respectively, we see that the degradation due to mismatch is eliminated. In this case, however, the conventional time-varying equalizer with 5 taps for each sub-ADC also works well, matching the performance of the time-frequency multiplier approach. This contrasts with the performance of the (symbol rate sampled) conventional scheme in Figure 11(a), and indicates that an *oversampled* conventional scheme is competitive with the time-frequency multiplier scheme. The time-frequency multiplier requires an FFT of twice the length, but has a smaller number of taps. In our example, the conventional scheme requires $5L$ taps, five each for the L time-varying filters corresponding to each sub-ADC, with the number of taps depending on the desired resolution. For the time-frequency multiplier approach, we need $2L$ taps for the two multipliers: $u_1[m]$ and $u_2[m]$ (periodic with period L from [13]), and 3 taps for the filters: $V_1[m]$ and $V_2[m]$, with a total of $2L + 3$ taps. The number of coefficients is insensitive to the resolution, unlike the conventional approach. Further research is required to quantify the tradeoffs in choosing between these approaches, but it is clear that oversampling is important in both settings.

V. Conclusions

We have presented promising preliminary results for two different, and complementary, approaches to addressing the ADC bottleneck in realizing mostly digital multi-Gigabit transceiver architectures. Many issues need further exploration, however. For communication systems using low-precision ADC, algorithms for timing and carrier synchronization, along with transmit precoding strategies, need to be devised and evaluated. For systems based on TI-ADCs, we must explore power and complexity tradeoffs in estimating and compensating for mismatch in a number of scenarios, including OFDM, single-carrier and MIMO systems.

References

[1] R. Walden, *Analog-to-Digital Converter survey and analysis*, IEEE J. Select. Areas Comm., 17(4):539–550, Apr. 1999.

[2] K. H. Lundbery, *High-speed analog-to-digital converter survey*, avaiblable at http://web.mit.edu/klund/www/papers/UNP_flash.pdf, 2002.

[3] 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.

[4] J. Singh, O. Dabeer, U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *to appear in IEEE Trans. Comm.*

[5] 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.

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

[7] D. Divsalar and M. K. Simon, *Multiple-symbol differential detection of MPSK*, IEEE Trans. on Comm., vol. 38, pp. 300-308, Mar. 1990.

[8] D. Warrier and U. Madhow, *Spectrally efficient noncoherent communication*, IEEE Trans. Info. Theory, vol. 48, pp. 651-668, Mar. 2002.

[9] J. Elbornsson, F. Gustafsson and J. E. Eklund, *"Analysis of Mismatch Effects in a Randomly Interleaved A/D Converter System"* , IEEE Trans. Circuits and Sys., vol. 52, Mar. 2005.

[10] M. Seo, MJW Rodwell and U. Madhow, *"Comprehensive digital correction of mismatch errors for a 400-Msamples/s 80-dB SFDR time-interleaved analog-to-digital converter"* , IEEE Trans. Microwave Theory and Techniques, vol. 53, pp. 1072-1082, March 2005.

[11] U. Madhow, *"Fundamentals of Digital Communication"* , Cambridge, 2008

[12] P. Sandeep, U. Madhow, M.Seo, M. Rodwell *"Joint Channel and Mismatch Correction for OFDM Reception with Time-interleaved ADCs: Towards Mostly Digital MultiGigabit Transceiver Architectures"* , IEEE GLOBECOM, Nov 2008, New Orleans, USA.

[13] P. Sandeep and U. Madhow, *"Scalable Mismatch Correction for Time-interleaved Analog-to-Digital Converters in OFDM Reception"* , technical report available at http://www.ece.ucsb.edu/wcsl/publications.html.

[14] S. Huang and B. C. Levy, *"Adaptive Blind Calibration of Timing Offset and Gain Mismatch for Two-Channel Time-Interleaved ADCs"* , IEEE tran. Circuits and Systems- I, vol. 53, pp. 1278-1288, June 2006.

[15] J. Elbornsson, F. Gustafsson, and J. E. Eklund, *"Blind Equalization of Time Errors in a Time-Interleaved ADC System"* , IEEE tran. Signal Processing, vol. 53, pp. 1413-1424, April 2005.

[16] T. Strohmer and J. Tanner, *"Fast Reconstruction methods for Bandlimited functions from Periodic Nonuniform Sampling"* , SIAM J. Numerical Analysis, vol. 44(3), pp 1073-1094, 2006.