STATISTICS GATHERING CONVERTERS: SYSTEM LEVEL METRICS, SIMULATED PERFORMANCE, AND PROCESS VARIATION ROBUSTNESS

Andrew J. Bean and Andrew C. Singer*

Department of Electrical and Computer Engineering University of Illinois at Urbana-Champaign Urbana, IL 61801, USA Email: {ajbean, acsinger}@illinois.edu

ABSTRACT

Analog to digital conversion is often a critical component of a digital communication link. However, the designs of typical architectures for analog to digital converters (ADCs) are focused primarily on signal reconstruction rather than gathering information for the reliable detection of symbols sent through a channel. Therefore, we consider new architectures for statistics gathering converters (SGCs), and demonstrate that these architectures achieve good communication performance while removing the artificial constraints imposed by the typical ADC design metrics.

In this paper, we extend previous work on system level metrics for statistics gathering converters (SGCs). For the particular case of the delay-line based SGC, we demonstrate two important facts. First, we consider the comparison between the performance indicated by system level metrics (BER and LMMSE) with the results of a simulated communication scenario utilizing a low complexity least mean squares equalizer. Simulations demonstrate that the system level metrics are an accurate representation of the realizable communication performance of a system using the converter in question, and that such performance can be nearly achieved by the delay-line SGC using a specially designed low complexity (LMS) equalizer that takes into account the particular structure of the SGC. Second, we demonstrate that the communication performance of the delay-line SGC is robust to significant levels of process variation, which manifest in random realizations of the values of the various SGC circuit elements. Notably, this is contrary to the strict requirements on process variation imposed by traditional metrics (SNDR, SFDR, THD) on conventional analog to digital converter designs.

Index Terms— ADC, LMMSE, BER, communication, system level metric

1. INTRODUCTION

Digital communication is the process of sending digital data, in the form of bits, from one location, the transmitter, to another, the receiver. Regardless of the origins of the data, the goal of the communication link is to reliably receive all of the data that was sent, or to do so with a bit error probability (BER) that is as small as possible. To achieve this task, the process of inserting redundancy into the data, in the form of forward error correction, converting the binary digital data into a waveform for transmission, and subsequent detection and estimation of the transmitted data is undertaken focussing on this system-relevant metric of performance, namely the bit error rate of the link. While the use of a system-relevant metric for link-level algorithm and architecetural designs makes sense, many of the critical components in such designs including circuit and system components are often designed using waveform-centric metrics, such as signal to noise plus distortion ratio, or the total harmonic distortion [1], which consider distortion caused to a sinusoidal input, when the input is reconstructed from its acquired samples, in the case of an analog-to-digital converter.

Initial work using analog-to-digital converter (ADC)-based receivers for 10Gb/s wireline and optical transceivers leveraged the power of DSP-based back-ends making use of modest resolution ADCs [2, 3] for subsequent data detection. As rates scale and resolution becomes more challenging, digital calibration has been increasingly employed [4], again focussing on minimizing converter nonlinearities due to ladder offsets or gain and phase mismatches in time-interleaved ADCs [5]. Rather than calibrating the ADC to improve such waveform-centric metrics, [1] considered a flash converter structure in which the sampling and reconstruction levels for the ADC are adjusted to minimize the link BER, resulting in dramatically improved link BER performance for ISI-dominated links, in the low resolution regime typical for the 10Gb/s-100Gb/s range. The information-theoretic capacity (maximum achievable rate with vanishing error probability) of a digital communication link comprising an additive white Gaussian noise channel followed by a lowresolution quantizer was studied in [6-8], along with strategies for reducing converter resolution while maintaining link performance. In [9], the information-theoretic concept of mutual information (MI) was used to explore the relationship between sampling-phase and maximum achievable data rate through a time-interleaved ADC architecture, in which the relative timing phase between branches of the ADC may not be uniform. In this work, it was demonstrated that even though waveform-centric metrics, such as THD or SNDR, degrade when such sampling phases are non-uniform, the achievable rate of a communication link is relatively insensitive to, and can even improve, through the use of non-uniform sampling across the ADC. In [10], mutual information was used to guide the design of achievable-rate optimal nonuniform quantizers for communication links, again demonstrating that, while SNDR and THD would degrade, achievable rate optimal designs have markedly non-uniform comparator thresholds. In [11], we used mutual information, linear minimum mean square error, and bit error rate as system level metrics to evaluate the quality of delay-line based "statistics gathering converter," in contrast with traditional ADCs, for communications applications.

One approach to simplifying the design of high-speed converter

^{*}This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA.



Fig. 1. A transmitter and channel connected to an SGC front-end, consisting of a delay line with five taps that may connect to samplers.

structures is to abandon the goal of analog-to-digital conversion altogether, and replace it with the true system-level goal of the analog front-end in a communication link; namely, to acquire statistics from the received signal waveforms that are sufficient for the problem of detecting the data that was transmitted. As in [11], rather than focussing on whether or not the waveform with all of its temporal properties are preserved by sampling, we focus on the development of a "Statistics Gathering Converter" or (SGC) whose primary role is to gather statistics for subsequent processing that will attempt to recover the transmitted data with low error probability.

Some of the challenges in the design of time-interleaved ADCs is maintaining constant gain and sampling phase across the branches, requiring considerable calibration and processing circuitry [2]. Similar challenges arise in flash architectures, maintaining uniformly increasing offsets across comparator ladders, while maintaining uniform gain and bandwidth characteristics, again, focussing calibration on waveform-centric metrics. However, for a digital communication link, valuable resources (such as power or chip area) might be better spent minimizing the overall link bit error rate, or maximizing the information capacity of the link.

Following up on our work on SGCs from [11], the goal of this paper is to continue moving away from ADCs, which have become the most power-hungry, sensitive component in the front-end of a communication link, due to their focussing on preserving irrelevant or unnecessary waveform-centric metrics. In their place, we will consider more general architectures that can be considerably less complex, require substantially lower power for operation, and be made less sensitive to circuit nonidealities, by focussing on simply preserving the information content of the gathered statistics, rather than maintaining waveform integrity. Specifically, we follow up the analytic evaluation of the delay-line SGC proposed in [11], shown in Figure 1, using extensive simulations to demonstrate two particular facts: First, that the system level metrics are an accurate representation of the realizable communication performance of a system using the converter in question, and that such performance can be nearly achieved by the delay-line SGC using a specially designed low complexity (LMS) equalizer that takes into account the particular structure of the SGC. And second, that the communication performance of the delay-line SGC is robust to significant levels of process variation, which manifest in random realizations of the values of the various SGC circuit elements. This is notably in stark contrast with the strict requirements on process variation imposed by traditional metrics (SNDR, SFDR, THD) on conventional analog to digital converter designs.

The organization of the paper is as follows: We review the system model in Section 2. We then review the system level metrics in Section 3. Finally, we provide simulation results and final remarks in Section 4.



Fig. 2. Basic model of a communication system.

2. COMMUNICATION SYSTEM MODEL

For this work, we use the system model specified in [11], which we briefly review in this section. For a more detailed description of our model, we refer the reader to [11].¹ We model the communication system as in Figure 2. This is a simple digital communication link with a pulse amplitude modulation (PAM) transmitter and a data converter front-end to the receiver. The sequence of transmitted symbols, $x[n] \in \mathbb{R}$, is modulated onto the channel through the D/C converter, such that the transmitter output is

$$x_1(t) = \sum_{i=-\infty}^{\infty} x[i]\tilde{p}(t-iT).$$
(1)

Hence, we transmit at a rate of one symbol every T seconds, or 1/T symbols per second. Without loss of generality, the dispersive effects of the communication channel are lumped within the modulator pulse shape $\tilde{p}(t)$. Additive channel noise is modeled by the zero mean wide sense stationary Gaussian process $w_c(t)$. The signal $x_2(t) = x_1(t) + w_c(t)$ is received by the converter at the front-end of the receiver.

A typical ADC front-end converter is a simple periodic sampler, such that the converter output samples are $y[n] = x_2(nS) + w_{\text{th}}(nS)$, where the converter samples the input signal every S seconds, and incurs some noise. This noise may consist of thermal noise, aperture jitter, or comparator ambiguity, which we lump into a single noise term [12]. We assume that each sample $w_{\text{th}}(nS)$ is independent of the others and zero mean IID Gaussian distributed.

In this work, we will compare the standard front-end with the particular SGC shown in Figure 1, which we will refer to as a delayline SGC. This consists of an amplifier feeding a simple passive network. Five outputs are labeled, from which samples may be taken every S seconds. Similarly with the ADC front-end noise, we assume that there may be noise injected in the SGC in addition to the channel noise.

For both the ADC and SGC front-ends, note that a finite input block of N symbols x[n] will produce a finite output block y[n].

¹For a copy of [11], please refer to IEEE Xplore or email the authors.



Fig. 3. LMMSE Metric and measured MSE vs. sync time.



Fig. 4. BER Metric and measured BER vs. sync time.

Furthermore, each of these systems is linear as a whole from input block to the block of observations, which motivates us to derive a linear discrete time model for the systems, as follows:

$$\mathbf{y}_{(N)} = \mathbf{A}_{(N)}\mathbf{x}_{(N)} + \mathbf{v}_{(N)},$$

where we explicitly denote the dependence on the input block size N. We have that $\mathbf{x}_{(N)}$ is a vector of all input symbols; $\mathbf{y}_{(N)}$ is a vector of all output observations from the samplers; $\mathbf{A}_{(N)}$ is a matrix representing the composite effects of the modulation pulse shape, the channel, and the front end; and $\mathbf{v}_{(N)}$ is a composite noise term that includes both the noise introduced by the channel and the noise introduced within the front-end circuitry. Due to space limitations, we refer the reader to [11] for a detailed derivation of the particular form of $\mathbf{A}_{(N)}$ and the statistical characterization of $\mathbf{v}_{(N)}$.

3. SYSTEM LEVEL METRICS - REVIEW

In [11], we proposed three system level metrics for evaluating data converters (e.g. SGCs and ADCs) used in communication applications. These metrics are based on mutual information (MI), linear minimum mean square error (LMMSE), and bit error rate (BER). In our results, we compare measured simulated performance with the LMMSE and BER metrics, which we briefly review here.

The LMMSE metric is given by the following:

$$\begin{split} D(\mathbf{A}, \mathbf{v}) &= \\ \lim_{N \to \infty} \frac{1}{N} \mathrm{tr} \left(\mathbf{R}_{\mathbf{x}_{(N)}} - \mathbf{R}_{\mathbf{x}_{(N)}} \mathbf{A}_{(N)}^{H} \mathbf{R}_{\mathbf{y}_{(N)}}^{-1} \mathbf{A}_{(N)} \mathbf{R}_{\mathbf{x}_{(N)}} \right), \end{split}$$

where we have defined

$$\mathbf{R}_{\mathbf{y}_{(N)}} = \mathbf{A}_{(N)} \mathbf{R}_{\mathbf{x}_{(N)}} \mathbf{A}_{(N)}^{H} + \mathbf{R}_{\mathbf{v}_{(N)}},$$

and the notation $tr(\cdot)$ indicates the matrix trace. The matrices $\mathbf{R}_{\mathbf{x}(N)}$, $\mathbf{R}_{\mathbf{y}(N)}$, and $\mathbf{R}_{\mathbf{v}(N)}$ are the covariance matrices of the vectors $\mathbf{x}(N)$, $\mathbf{y}(N)$, and $\mathbf{v}(N)$, respectively. We will define $\mathbf{R}_{\mathbf{x}(N)} = \mathbf{I}$, so it is noted that the expression may be simplified. This expression can be used for arbitrary zero mean unit variance IID input distributions.



Fig. 5. MSE of realizations of SGCs under process variation.



Fig. 6. BER of realizations of SGCs under process variation.

The objective of design based on LMMSE would be to minimize the value of $D(\mathbf{A}, \mathbf{v})$. Note that this metric assumes a linear equalizer of unbounded complexity. In this work, we demonstrate that a low complexity equalizer, constructed to take advantage of the particular SGC's structure, is sufficient to achieve the performance indicated by this metric.

The other metric we will use is bit error rate (BER), which will be denoted $p_e(\mathbf{A}, \mathbf{v})$. The assumption when we use this metric is that the input symbols x[n] are chosen from the finite alphabet $\{-1, +1\}$ with equal probability. Hence, we again have a zero mean unit variance input distribution. We define the BER as the average probability of detection error from the LMMSE estimate of the input symbols. Let $Q(x) = \Pr[X > x]$, where X is a zero mean unit variance Gaussian random variable. Let $\hat{\mathbf{x}}_{(N)}$ be the LMMSE estimate of $\mathbf{x}_{(N)}$. Define $\mathbf{H}_{(N)} = \mathbf{A}_{(N)}^H \mathbf{R}_{\mathbf{y}_{(N)}}^{-1}$. Then we have that $E[\hat{\mathbf{x}}_{(N)}|\mathbf{x}_{(N)}] = \mathbf{H}_{(N)}\mathbf{A}_{(N)}\mathbf{x}_{(N)}$ and $\mathbf{R}_{\hat{\mathbf{x}}_{(N)}}|\mathbf{x}_{(N)} =$ $\mathbf{H}_{(N)}\mathbf{R}_{\mathbf{v}_{(N)}}\mathbf{H}_{(N)}^H$. Finally, this gives us the BER metric as

$$p_e(\mathbf{A}, \mathbf{v}) = \lim_{N \to \infty} E_{\mathbf{x}_{(N)}} \left[\frac{1}{N} \sum_{i=1}^N Q\left(\frac{x[i]E[\hat{x}[i]|\mathbf{x}_{(N)}]}{\sigma_{(N)}[i]} \right) \right],$$

where $(\sigma_{(N)}[i])^2$ are the diagonal elements of $\mathbf{R}_{\mathbf{\hat{x}}_{(N)}|\mathbf{x}_{(N)}}$. The objective of design based on bit error rate would obviously be to minimize the value of $p_e(\mathbf{A}, \mathbf{v})$. Again, we demonstrate that a low complexity equalizer that takes advantage of the SGC's structure is sufficient to achieve the performance indicated by this metric.

4. SIMULATION RESULTS

The particular communication system we will consider in this paper utilizes a delay-line SGC front-end architecture, and is shown in Figure 1. Here, we have a transmitter that generates the voltage signal $x_1(t)$. This is connected to a channel consisting of an RC low pass filter with $R_{ch}C_{ch} = 0.1$ sec. The noise signal $w_c(t)$ is added to the output of the channel to produce $x_2(t)$, which is the input to the front-end. The front-end has an input buffer to isolate the dynamics of the passive delay-line from that of the channel. This buffer feeds a chain of resistors, inductors, and capacitors. The nominal values for these components are chosen to be $R_1 = R_2 = 0.01\Omega$, L = 0.7H, and C = 1.0F, such that the delay per section is approximately 1sec. This can be thought of as normalized time. An additional terminating resistor is at the end of the chain with $R_{end} = 1.0\Omega$. For all of our simulations, the observations used for the recovery of the transmitted data come from sampling outputs $y_1(t)$, $y_3(t)$, and $y_5(t)$ in Figure 1. Symbols are transmitted at 1 symbol every 2 seconds, and each of the chosen outputs is sampled simultaneously once every 6 seconds, for an average of one observation per input symbol. Symbols $x[i] \in \{-1, +1\}$ modulate a square pulse p(t) = u(t) - u(t - 2), which is then filtered by the described channel resulting in $\tilde{p}(t)$ in Equation (1).

We apply our assumptions from above about the input distribution and noise component within the front-end. The channel noise has the following autocorrelation function:

$$R_{w_{\mathrm{c}}}(t) = \sigma_{\mathrm{ch}}^2 \exp\left(rac{-|t|}{R_{\mathrm{ch}}C_{\mathrm{ch}}}
ight).$$

In our simulations, we compare statistics gathering front-ends to traditional ADC front-ends, which are simply modeled as sampling the signal $y_1(t)$ in Figure 1. (Note that the delay-line has no effect on the signal $y_1(t)$.) For fair comparisons between the different front-ends, we fix the analog channel noise autocorrelation and the discrete time thermal (front-end) noise variance, and define the SNR as

$$SNR = \frac{E[x[i]^2]}{R_{w_c}(0) + \sigma_{th}^2} = \frac{1}{\sigma_{ch}^2 + \sigma_{th}^2}.$$
 (2)

As with the SGC, we configure the ADC to generate, on average, one observation per input symbol. Hence, S = 2sec for the ADC.

Accuracy of signal reconstruction is of no fundamental importance when the objective of data conversion is the eventual recovery of transmitted data. As compared with converters designed for high reconstruction accuracy, converter architectures that allow the introduction of signal distortions can still preserve relevant statistics about the transmitted signal for the purpose of data recover. Furthermore, they can actually improve the performance in this task while simultaneously removing the need for precise fabrication and calibration. It is the goal of this section to demonstrate this by comparing a system using the described delay line SGC with a comparable design utilizing a classical ADC.

We first examine the relevance of the LMMSE and BER system level metrics for the design of practical systems. This is important because the metrics assume linear detectors of unbounded complexity, wherease applications require finite, and often low, complexity. We compare the performance indicated by the system metrics with that achieved by carefully designed low complexity equalizers. Specifically, the SGC equalizer consists of three filters, each with 11 taps. Let $\{y_1[3i]\}, \{y_3[3i-1]\}$ and $\{y_5[3i-2]\}$ for integers i be the sequences of outputs from the respective delay line taps 1, 3, and 5. Then the full interleaved observation sequence is $\{..., y[0], y[1], ...\} = \{..., y_1[0], y_5[1], y_3[2], y_1[3], y_5[4], ...\}$. To estimate symbol x[i], we use filter $h_{mod(i,3)}$ to get the estimate $\hat{x}[i] =$ $h_{\text{mod}(i,3)}^{\top} \mathbf{y}[i]$, where $\mathbf{y}[i] = (y[i-5], ..., y[i+5])^{\top}$. In this work, we use a simple LMS update with a step size parameter μ to independently adapt each of the estimation filters. The need for the three separate estimation filters arises because the characteristics of the observation vector $\mathbf{y}[i]$ are similar to those of $\mathbf{y}[i+3j]$ for integer values j, but potentially quite different from $\mathbf{y}[i+3j+1]$ and $\mathbf{y}[i+3j+2]$. This can be seen in the channel matrix **A**. For a baudsampled system using an ADC, A will have a toeplitz structure, but

in a system using our SGC (with the outputs arranged in a vector as $\{y[i]\}$ is here), it will instead have a *blockwise* toeplitz structure. The width of the blocks composing **A** determine the number of separate estimation filters to maintain and adapt.

Figures 3 and 4 show how our system level metrics compare with the actual performance of a communication system utilizing either an SGC or ADC with postproccessing of the observations by appropriate low complexity equalizers. On the horizontal axis, we have the synchronization time of the samplers (both ADC and SGC), with the center of the eye being at odd integer times 2i + 1. The thick lines in both figures are for the SGC, thin for the ADC. The dashed lines are for the value of the system level metric achieved for the particular synchronization time, whereas the solid lines are for the measured MSE and BER of a communications utilizing the respective data converter with the appropriate equalization method. For measuring the MSE and BER of the SGC systems, we trained the length 11 specialized LMS equalizer with parameter $\mu = 0.01$ on 30000 symbols and measure the MSE and BER on 2e7 subsequent symbols. For ADC systems, we do the same, but with a standard LMS equalizer. The figures show that the measured performance achieved by systems with constrained complexity is quite similar to the theoretical performance indicated by the system level metrics. In particular, the system level metrics correctly indicate that a system utilizing an SGC has the potential to outperform one with an ADC. Furthermore, the system level metrics and the corresponding measured performance have the same qualitative characteristics.

In the next set of simulations, we show that the communications performance of the SGC front-end, when paired with the described low complexity adaptive equalizer, is still able to beat the performance of the ideal ADC under significant levels of deviation from the nominal circuit parameters. In particular, recall that we have four each of capacitors and inductors in the SGC with nominal values L = 0.7H and C = 1.0F. We model process variation in the circuit by letting the true circuit parameter be Gaussian distributed with the nominal value as mean, and standard deviation of 0.05 (H or F). Figures 5 and 6 show the results of 2800 random instances of the SGC system, where we train the length 11 specialized LMS equalizer with parameter $\mu = 0.01$ on 30000 symbols and measure the MSE and BER on 2e7 subsequent symbols. Each point in the plots represents the measured MSE / BER versus the norm of the deviation from the nominal parameters, i.e. $|P - \bar{P}|_2$, where P is the vector of 8 circuit parameters and \overline{P} is the vector of 8 nominal values. The dashed lines indicate the performance level of the system when using a classical ADC. It is clear that the performance is robust to the parameter variation, and remains better than the ADC in all cases. Similar results hold when the standard deviation of the parameters is doubled to 0.1, with a small number of instances (5 out of 1000 points for both MSE and BER) with high deviation from the nominal parameters having worse performance than the ADC.

We have proposed to throw out the standard notions of quality that guide most data converter designs and replace them with metrics that are more relevant to our goal of digital communication. In support of this view, we have demonstrated that a relaxed specification data converter, or "SGC", is in fact competitive with and potentially better than an ADC that is "ideal" with respect to the conventional metrics. These metrics remain relevant even for the design of low complexity systems, and they help elucidate wiggle room in the circuit specifications where the standard metrics indicate that no such room exists. Much further work remains, including study of quantization effects, further analysis of process variation sensitivity, completely different SGC architectures, and figures of merit to further facilitate comparisons between converter designs.

5. REFERENCES

- M. Lu, N.R. Shanbhag, and A.C. Singer, "BER-optimal analog-to-digital converters for communication links," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, 2010, pp. 1029–1032.
- [2] A. Nazemi et al., "A 10.3GS/s 6bit (5.1 ENOB at nyquist) timeinterleaved/pipelined ADC using open-loop amplifiers and digital calibration in 90nm CMOS," in 2008 Symposium on VLSI Circuits, Technical Digest of Papers, 2008, pp. 18–19.
- [3] H. M. Bae, J. Ashbrook, J. Park, N. Shanbhag, A.C. Singer, and S. Chopra, "An MLSE receiver for electronic-dispersion compensation of OC-192 links," *Journal of Solid-State Circuits*, vol. 41, no. 11, pp. 2541–2554, November 2006.
- [4] P. Nikaeen and B. Murmann, "Digital compensation of dynamic acquisition errors at the front-end of high-performance A/D converters," *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 3, pp. 499–508, 2009.
- [5] C. Grace, P. J. Hurst, and S. H. Lewis, "A 12b 80MS/s pipelined ADC with bootstrapped digital calibration," in *IEEE International Solid-State Circuits Conference (ISSCC) Digest* of Technical Papers, February 2004, pp. 460–461.
- [6] J. Singh, O. Dabeer, and U. Madhow, "Communication limits with low precision analog-to-digital conversion at the receiver," in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2007.

- [7] J. Singh, O. Dabeer, and U. Madhow, "Capacity of the discretetime AWGN channel under output quantization," in *Proceed*ings of the IEEE International Symposium on Information Theory (ISIT), 2008.
- [8] O. Dabeer, J. Singh, and U. Madhow, "On the limits of communication performance with one-bit analog-to-digital conversion," in *Proceedings of the IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2006.
- [9] A.C. Singer, A. J. Bean, and J. W. Choi, "Mutual information and time-interleaved analog-to-digital conversion," in *Information Theory and Applications Workshop*, 2010, pp. 1–5.
- [10] G. Zeitler, A. Singer, and G. Kramer, "Low-precision A/D conversion for maximum information rate in channels with memory," *IEEE Transactions on Communications*, vol. 60, no. 9, pp. 2511–2521, September 2012.
- [11] Andrew J. Bean and Andrew C. Singer, "The SGC: A simple architecture for gathering statistics in communication links," in *IEEE Workshop on Signal Processing Systems*, 2013.
- [12] R. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, April 1999.