System-assisted Analog Mixed-signal Design

Naresh Shanbhag and Andrew Singer
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, Urbana, IL-61801, USA
Email: [shanbhag,acsinger]@illinois.edu

***Abstract*—In this paper, we propose a system-assisted analog mixed-signal (SAMS) design paradigm whereby the mixed-signal components of a system are designed in an application-aware manner in order to minimize power and enhance robustness in nanoscale process technologies. In a SAMS-based communication link, the digital and analog blocks from the output of the information source at the transmitter to the input of the decision device in the receiver are treated as part of the composite channel. This comprehensive systems-level view enables us to compensate for impairments of not just the physical communication channel but also the intervening circuit blocks, most notably the analog/mixed-signal blocks. This is in stark contrast to what is done today, which is to treat the analog components in the transmitter and the analog front-end at the receiver as transparent waveform preservers. The benefits of the proposed system-aware mixed-signal design approach are illustrated in the context of analog-to-digital converters (ADCs) for high-speed links. CAD challenges that arise in designing system-assisted mixed-signal circuits are also described.**
[1]

I. Introduction

Moore's Law, the driving force behind the global semiconductor industry for the last 50 years, is under threat today from artifacts of nanoscale dimensions. Process, voltage and temperature (PVT) variations, leakage, soft errors, and noise in sub-45nm process technologies [4] are conspiring to make it difficult to reap the benefits of feature size scaling due to reliability concerns in both digital and analog mixed-signal (AMS) circuits and systems. A parallel trend is the growing functional complexity and power requirements of next generation applications. The result is a power and reliability problem in nanoscale systems-on-a-chip (SOCs). Reliability and power are interlinked problems viewed by the semiconductor industry as the key inhibiters of Moore's Law. Not surprisingly, since 2001, the International Technology Roadmap for Semiconductors (ITRS) [1] has stated the achievement of reliability and energy-efficiency as two of the grand challenges facing the semiconductor industry.

Various approaches are presently being explored for designing robust and energy-efficient digital circuits and systems. These include noise and variation-tolerant circuits, architectures and systems. We proposed a *communications-inspired* design paradigm [6], whereby the nanoscale circuit fabric is viewed as a noisy communication channel or a network, and statistical design techniques are employed to compensate for hardware errors due to device and circuit imperfections. More recently, stochastic computation [11] has been proposed as a systematic approach for designing communications-inspired computing platforms. Stochastic computing exploits the statistical nature of application-level performance metrics of emerging applications, and matches it to the statistical attributes of the underlying device and circuit fabrics. The benefits of such a design philosophy are the tremendous gains in robustness and energy-efficiency in presence of an extremely high degree of unreliability (e.g., error-rates of up to 90%) in the circuit fabric. For example, orders-of-magnitude enhancement in system reliability have already been demonstrated for filtering [6], motion-estimation, Viterbi decoding, and CDMA pseudo-noise (PN) acquisition kernels [13], among others, along with 30%-to-60% energy-savings.

Moore's Law scaling is designed to favor digital systems. Today, in spite of the multitude of challenges facing nanoscale digital circuits and systems, it is AMS systems that prove to be the bottleneck when it comes to scaling. Therefore, it is important to consider non-traditional approaches to AMS design in nanoscale process technologies. Analog mixed-signal circuits and systems, though a critical part of any modern day embedded applications, derive limited benefits from process scaling. The success of the communications-inspired design philosophy in computing motivates us to explore its application in the AMS domain, which we refer to as *system-assisted analog mixed-signal* (SAMS). Of particular interest is the role of AMS components in a high-speed, e.g., backplane or optical communication link. This paper outlines key challenges in the design of SAMS-based AMS blocks in such links, with concrete examples.

This paper is organized as follows: Section II describes the system-assisted design paradigm and contrasts it with the prevailing approach, which is commonly referred to as digitally-assisted analog. In section III, we illustrate the application of the SAMS design approach to low resolution, high-speed analog-to-digital converters (ADCs). Section IV describes challenges in modeling SAMS components, and concludes by outlining CAD challenges in implementing SAMS systems.

II. System-assisted Mixed-Signal (SAMS) Design

Consider a conventional communication link as shown in Fig. 1(a), where reliable data transfer is achieved over a noisy channel via coding, whereby the encoder imposes specific structure on the transmitted data. The receiver exploits this structure to recover the data in the presence of noise and other channel impairments. This post-Shannon era link, however,

[1]978-3-9810801-7-9/DATE11/©2011 EDAA

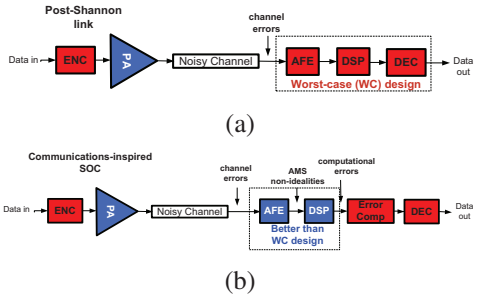

Fig. 1. Communications-inspired system-assisted AMS (SAMS) design: (a) conventional link, and (b) a SAMS-based link.

assumes that the processing at the transmitter and the receiver is ideal or close to ideal, i.e., error free. As a result, the processing components (both digital and AMS) at the transmitter and the receiver are designed to meet the worst case process, voltage, temperature (PVT), and channel conditions. Worst case designs, though reliable, are extremely power hungry.

In contrast, from a communications-inspired SAMS viewpoint, the AMS blocks (both analog and digital) are treated as an extension of the physical channel as shown in Fig. 1(b). In doing so, we permit these blocks to exhibit a greater degree of non-ideality than permitted in conventional systems, and exploit application-level awareness to compensate for them. In other words, a SAMS-based communications link is designed by treating all the blocks from the output of the information source at the transmitter to the input of the decision device in the receiver as part of a composite unreliable channel. This comprehensive view of the channel enables us to compensate for impairments of not just the physical channel but also the intervening circuit fabric, most notably the AMS blocks. This is just the opposite of what is done today, which is to treat the analog front-end at the transmitter and the receiver as transparent waveform preservers, i.e., meeting a *fidelity criterion*. Such a fidelity criterion strives to preserve the analog waveform shape, and in doing so certainly captures the information resident in the received signal and a whole lot more. However, such a fidelity criterion ignores the true role of the analog front-end in the detection process that is so fundamental to a communication link, and in doing so results in overdesigned, power-hungry solutions. Furthermore, designs using a fidelity criterion are unsustainable in nanoscale CMOS process technologies as the latter do not favor AMS circuits. In contrast, by focusing on the ultimate goal of achieving reliable information transfer, we design the AMS blocks as part of the overall system goal to meet a *detection criterion*. Specifications derived from a system-level detection criterion can be much more relaxed than those corresponding to the waveform fidelity criterion and thus significant power savings can be expected.

The SAMS approach is distinct from the prevailing *digitally assisted analog* (DAA) approach [8] in that, though a commendable trend, the latter remains a component level technique. In DAA, a number of digital signal processing and other digital heavy techniques are employed to compensate for AMS non-idealities. What limits it to the component level is the use of the very same metrics e.g., spurious-free dynamic range (SFDR) and signal-to-noise-plus-distortion ratio (SFDR) in case of ADCs, as is employed in conventional analog heavy designs. In contrast, SAMS employs system level metrics that are quite distinct, and hence result in unconventional AMS architectures.

In the proposed approach, one can design the analog front-end architectures and the DSP back-end to be intrinsically imperfect with the degree of imperfection being matched to the reliability specifications on the data being processed. For example, in links with forward error-control (FEC), the output of a equalizer for the channel impairments needs to achieve a typical bit error-rate (BER) of 10^{-3} to 10^{-5}, whereas the output of the FEC decoder brings the BER down to 10^{-12} or smaller. This orders-of-magnitude discrepancy in the BER specifications between the equalizer and decoder has only recently been exploited [9] to reduce power in heavily power-constrained links such as back-plane. In this paper, we will focus on ADCs, as these are among the more power hungry blocks in high-speed links. Today, researchers in industry and academia continue to explore high-speed ADC designs using the traditional fidelity criteria (SFDR and SFDR), where the analog front-end (AFE) in the receiver is designed to behave as a transparent information conduit. In such links, low-power ADCs are particularly difficult to design, and the effective number of bits (ENOB) usually does not exceed six [3], [5], [10]. This paper describes how system-level information can be incorporated into the design of information preserving, high-speed ADCs such that power is dramatically reduced without compromising the bit error-rate (BER) of the overall communication link.

III. BER-aware ADC Architectures

We investigate low-power, high-speed ADC architectures using a detection criterion whereby BER-optimal values of ADC parameters such as quantization/reference voltage-levels, quantizer resolution, and sampling phase are computed via signal processing techniques in order to preserve the transmitted information through the link, rather than to needlessly preserve the analog waveform properties. We expect at least an order-of-magnitude reduction in ADC power-consumption over a variety of wireless and wireline channels.

Figure 2(a) shows the measured eye-diagram at the input to the ADC in a typical single-mode fiber (SMF) link. Other high-speed channels exhibit similar features. Note that the intersymbol interference and noise results in a non-uniform signal probability distribution. This feature (among others) can be exploited to reduce the ADC resolution needed to achieve the required BER via a non-uniform assignment of voltage references. Second, the optimal sampling phase of such a signal will not necessarily be at the center of the eye as is usually assumed in conventional systems. Finally, one can view the ADC itself as a time-sampled band-limited

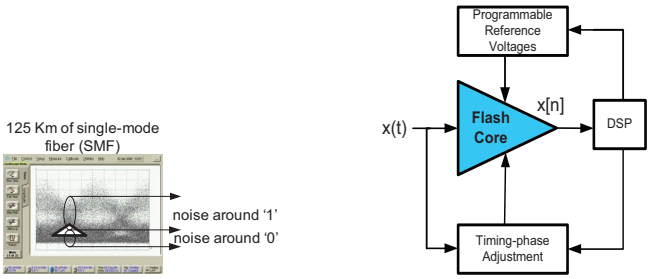

(a) (b)

Fig. 2. BER-aware ADC architectures: (a) eye diagram at the output of a $125km$ single-mode optical fiber, and (b) a general SAMS-based architecture.

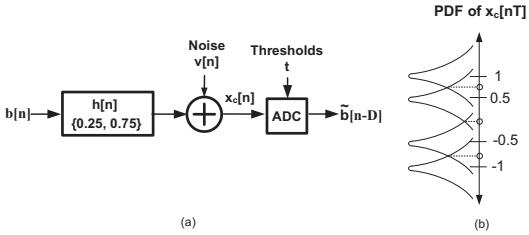

Fig. 3. Motivational example.

communication channel itself and thus digital signal processing techniques can be employed to compensate for its high-frequency/AC non-idealities. These steps will result in relaxed specifications for the ADC and hence opportunities for reduced power consumption.

The considerations mentioned above lead to the general SAMS-based high-speed ADC architecture shown in Fig. 2(b), where the digital signal processor (DSP) employs the discrete-time samples of the received signal to determine BER-optimal settings for the reference voltages, the sampling phase and the resolution (not shown). It is assumed that the ADC core comprises a flash converter though other architectures can also be considered.

A. Reference-level Optimized BER-aware ADC Architectures

We begin with a motivational example in order to illustrate the difference between a fidelity criterion such as $SNDR$ and a detection/information-based criterion such as BER. Consider a communication link (Fig. 3), where the transmitter employs 2-PAM modulation and the channel is represented by the discrete-time impulse response $[0.25\ 0.75]$. The receiver consists of a detector immediately following the ADC, i.e a decision on the transmitted symbols is made based on the ADC output. We apply Bayesian hypothesis testing with the ADC inputs $x_c[n]$ being treated as observations, since the ADC effectively acts as the detector.

$$\begin{aligned}\tilde{b}[n] = 1 \text{ if } \Pr\{b[n]=1|x_c[n]=y\} &> \Pr\{b[n]=1|x_c[n]=y\} \\ = \quad -1 \text{ else.} \end{aligned} \tag{1}$$

Now,

$$\Pr\{b[n]=1|x_c[n]=y\} = \frac{\Pr\{b[n]=1, x_c[n]=y\}}{\Pr\{x_c[n]=y\}} \tag{2}$$

$$= \frac{\Pr\{x_c[n]=y|b[n]=1\}\Pr\{b[n]=1\}}{\Pr\{x_c[n]=y\}} \tag{3}$$

Assuming that the transmitted symbols are independent and identically distributed,

$$\begin{aligned}\Pr\{x_c[n]=y|b[n]=1\} = \\ = \frac{1}{2}\sum_{b[n-1]=-1}^{1}\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\frac{(y-0.25-0.75b[n-1])^2}{2\sigma^2}} \\ = \frac{1}{2\sqrt{2\pi\sigma^2}}\left(e^{-\frac{(y-1)^2}{2\sigma^2}}+e^{-\frac{(y+0.5)^2}{2\sigma^2}}\right)\end{aligned} \tag{4}$$

Similarly,

$$\begin{aligned}\Pr\{b[n]=-1|x_c[n]=y\} = \\ = \frac{\Pr\{x_c[n]=y|b[n]=-1\}\Pr\{b[n]=-1\}}{\Pr\{x_c[n]=y\}}\end{aligned} \tag{5}$$

and

$$\Pr\{x_c[n]{=}y|b[n]{=}1\} = \frac{1}{2\sqrt{2\pi\sigma^2}}\left(e^{-\frac{(y+1)^2}{2\sigma^2}}{+}e^{-\frac{(y-0.5)^2}{2\sigma^2}}\right) \tag{6}$$

From Equations 1-6, the optimal detection rule can be given as,

$$\begin{aligned}\tilde{b}[n] &= 1 \text{ if } \frac{1}{2\sqrt{2\pi\sigma^2}}\left(e^{-\frac{(y-1)^2}{2\sigma^2}}+e^{-\frac{(y+0.5)^2}{2\sigma^2}}\right) \\ &> \frac{1}{2\sqrt{2\pi\sigma^2}}\left(e^{-\frac{(y+1)^2}{2\sigma^2}}+e^{-\frac{(y-0.5)^2}{2\sigma^2}}\right) \\ &= -1 \text{ else.}\end{aligned} \tag{7}$$

Solving the inequality (7), the observation $x_c[n]$ can be divided into four regions with three thresholds, and the decision rule is given by,

$$\tilde{b}[n] = 1 \text{ if } x_c[n] \in [-0.75,\ 0)\,[0.75,\ \infty) \tag{8}$$

$$= -1 \text{ else.} \tag{9}$$

This can be understood from noting that the noise-free channel output $z[n]$ takes four possible values from the set $[-1\ \ -0.5\ \ 0.5\ \ 1]$. The ADC input therefore, has a multi-modal distribution consisting of a mixture of Gaussian modes centered at each of these values, when there is additive Gaussian noise in the channel or due to thermal noise in the receiver. The modes centered at -0.5 and 1 correspond to a transmitted 1 and those centered at -1 and 0.5 correspond to a transmitted -1. The ADC thresholds are designed such that modes are mapped to the corresponding hypothesis on the transmitted symbol. The BER-optimal ADC in this case is clearly not SQNR optimal, as mapping the mode centered at -0.5 to a 1 and vice versa would incur a heavy SQNR cost. However, by employing the ADC for detection, rather than preserving sinusiodal waveforms (as implied by SNDR and SQNR criteria) no system-level performance is lost from a simple one-bit converter even in the presence of severe SNDR or SQNR penalty by traditional metrics.

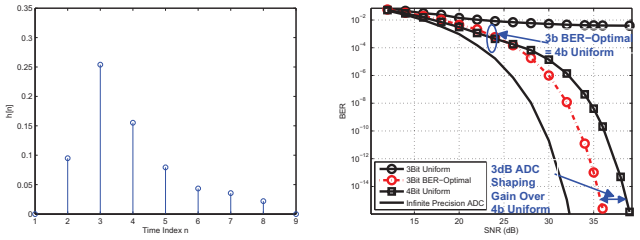

(a) (b)

Fig. 4. Performance of the reference-level optimized BER-aware ADC for an ISI channel with 2-PAM modulation and a LE: a) sampled impulse response, and b) BER vs. SNR curves for a 3-bit uniform, 3-bit BER-optimal, 4-bit uniform, and an infinite-precision ADC, respectively.

If we assume that the ADC output is processed by a linear equalizer (LE) or a decision-feedback equalizer (DFE), whose output is then sliced (quantized) in order to obtain the detected bits $\tilde{b}[n]$, then one can formulate a generalized detection criterion similar to (7) [7]. Such a criterion incorporates not just the channel parameters but also those of the subsequent equalizer, for an arbitrary number of reference levels and for arbitrary signal constellations. We define the *ADC shaping gain* as $S_G(BER) = SNR_{old}(BER)(dB) - SNR_{new}(BER)(dB)$ to quantify the reduction in SNR measured in dB achieved via the BER-optimal techniques.

Figure 4(a) shows the sampled impulse response for a channel with ISI, and Fig. 4(b) plots the BER vs. SNR for various ADC architectures. Note that the conventional 3-bit (uniform) ADC exhibits an error floor slightly below 10^{-2}, and thus isn't considered a viable option. However, a 3-bit BER-optimal ADC not only beats the 3-bit uniform ADC but also outperforms a 4-bit uniform ADC with an $S_G(BER = 10^{-15}) = 3$dB. The reduction in precision by 1-bit and the ADC shaping gain both result in power savings. Assuming a flash converter, just the reduction in precision itself provides approximately 50% reduction in power.

B. Sampling-phase Optimized BER-aware ADC Architectures

One approach to simplifying the design of high-speed converter structures is the use of time-interleaved architectures, where a rate $1/T$ converter is constructed by using M rate $1/MT$ converters operating in parallel, each operating on a different phase of an M-phase clock. While each of the sample-and-hold circuits within the M branches in such a design must still have the instantaneous bandwidth of the overall rate, i.e. $1/T$, the design can be relaxed by providing MT seconds for the conversion process of the samples. One of the key challenges in the design of such time-interleaved ADCs is maintaining constant gain and sampling phase across the branches, and as a result, considerable calibration and processing circuitry is employed in such architectures [2]. However, once again, the focus of such extended efforts in calibration are not on minimizing the BER of the communication link, but rather in maximizing the signal to noise plus distortion ratio (SNDR) or in minimizing the total harmonic

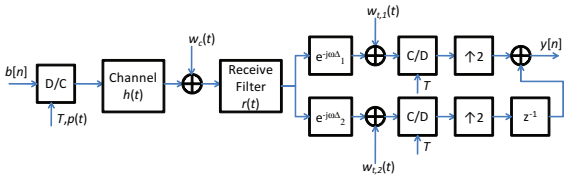

Fig. 5. Block Diagram of a digital communication link with a time-interleaved analog to digital converter in the receiver. When the C/D converters each operate at one sample every T seconds (as shown), the front-end is fractionally-spaced; when the converters operate at one sample every $2T$ seconds, the converter is symbol spaced.

distortion induced by such offsets. This amounts to assuming that the signals of primary interest to the receiver are sinusoidal signals, and that it is this distortion from a pure sinusoidal tone that is to be avoided. However, for a digital communication link, rather than spending valuable resources within the analog front-end preserving the quality of transmitted tones, these resources might be better spent in attempts to either minimize the overall link bit error rate, or to maximize the information capacity of the link. The goal of the paper [12], therefore was to begin to analyze the information capacity of such ADCs, with a particular focus on studying the mutual information between the transmitted bits through such a link and the sampled outputs of such a time-interleaved ADC.

The block diagram shown in Fig. 5 depicts a simple digital communication link with a PAM transmitter and a time-interleaved ADC front-end to the receiver. The sequence of transmitted symbols, $b[n]$ are modulated onto the channel as indicated through the D/C converter, such that the transmitter modulates the transmit pulse $p(t)$, as

$$x(t) = \sum_{n=-\infty}^{\infty} b[n]p(t - nT),$$

at a rate of one symbol every T seconds, or $1/T$ symbols per second. The dispersive effects of the communication channel are modeled by the equivalent baseband impulse response $h(t)$, and additive channel noise is modeled by the wide sense stationary process $w_c(t)$. At the receiver, the receive waveform is given by

$$s(t) = \int_{-\infty}^{\infty} x(\tau)h(t - \tau)d\tau + w_c(t),$$

which is then processed with the receive shaping filter with impulse response $r(t)$, which also serves to bandlimit (or attenuate the out of band components of) the additive channel noise $w_c(t)$. For example, in the absence of a dispersive channel response $h(t)$, the transmit and receive shaping filters might be selected as square-root raised cosine filters to provide an ISI-free link. The subsequent stages shown to the right in Fig. 5 depict a two-phase time-interleaved ADC. The received signal is passed through one converter along the top branch, which takes a single sample every T seconds at a relative delay of Δ_1. The lower branch includes a different delay of

Δ_2 seconds prior to sampling at a rate of one sample every T seconds. Within the branch k of the receiver, independent additive thermal noise is modeled as additive stationary white Gaussian noise, $w_{t,k}(t)$, out to the resolution bandwidth of the sample and hold circuitry. These samples are then interleaved to produce a net of two samples every T seconds, or a fractionally-spaced converter. If each of the converters were to operate at a single sample every $2T$ seconds, then the converter would be symbol-spaced, rather than fractionally spaced. This process could be further parallelized to include M converters, each operating at a rate of one sample every MT seconds.

We note that the system shown in Fig. 5 has a linear input-output relationship. More specifically, consider a finite vector of symbols $\mathbf{b} = (b[0], ..., b[N-1])^T$ and an observation vector $\mathbf{y} = (y[0], ..., y[L-1])$. It is possible to relate these input and output vectors according to the relation

$$\mathbf{y} = \mathbf{A}(\mathbf{D})\mathbf{b} + \mathbf{w}_{th} + \mathbf{w}_{ch}. \tag{10}$$

In (10), $\mathbf{A}(\mathbf{D})$ is a channel matrix that depends on the relative sampling time delays of the ADC branches, $\mathbf{D} = (\Delta_1, ..., \Delta_M)^T$, $\mathbf{w}_{th}$ is a vector of stationary independent Gaussian noise samples that arise from the bandlimited sampling (due to the integration time of the sample and hold circuitry) of $w_{t,j}(t)$, and $\mathbf{w}_{ch}$ is a vector of noise samples that arise from sampling the output of the receive shaping filter due to the channel noise,$w_c(t)$.

In order to gain some insight into the potential effects of different values of $\mathbf{D}$ on the effectiveness of the ADC samples for communication, we examine the input to output mutual information, per channel use, as a function of the delays $\mathbf{D}$. As such, we investigate the following:

$$C(\mathbf{D}) = \lim_{N\to\infty} \frac{1}{N} I(\mathbf{b}^N; \mathbf{y}^{L(N)}) \tag{11}$$

In equation (11), it is implicit that the observation vector $\mathbf{y}^{L(N)}$ depends on the particular choice of $\mathbf{D}$. For simplicity of analysis and computation, we assume that the distribution of input symbols is fixed such that each symbol is an independent Gaussian random variable with mean zero and unit variance. The mutual information between the transmitted symbols and the output of the converter at the receiver was used as a guide to measure the optimal relative sampling phases for symbol-spaced and fractionally-spaced receiver structures. It was observed that for symbol-spaced receivers, the conventional approach of equispaced samples coincides with the locations of maximum mutual information. However for oversampling receivers, the relative delays of the arms of the converter were not equispaced (as would have been required for a fidelity criterion such as SFDR or SNDR), but rather were a function of the relative amounts of channel noise and circuit noise present in the converter samples. Simulations of a link using a turbo-equalization based receiver confirmed that the location of the optimal sampling points observed via mutual information coincided with the sampling phases that achieved the minimum bit error rate.

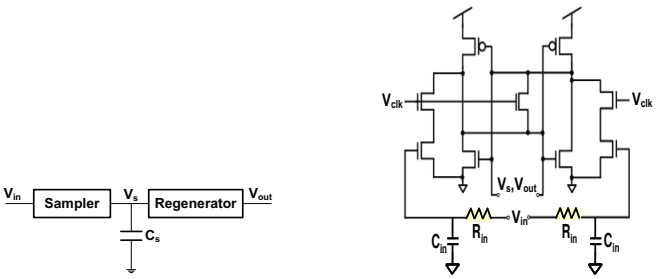

Fig. 6. The latch model: a) block diagram, and b) circuit schematic.

IV. CAD Challenges in SAMS

System-assisted AMS designs need behavioral component-level models. Unlike conventional design methodologies, SAMS requires system-level models for each component, whereby each component is characterized in terms of power and performance over a much broader range of operating conditions including ones where the component exhibits undesirable/erroneous behavior. This section illustrates some of these challenges in the design of a high-speed latch employed in ADCs.

A. Latch Modeling

A latch includes a sampler front-end and a regenerator as shown in Fig. 6(a), and employs a two-phase clock to sample the input signal V_{in} on a capacitor C_s during the *sampling* phase, and rails the output V_{out} in the *regenerate* phase. A high-speed static reset based digital latch shown in Fig. 6(b), is often employed to convert the analog output of a comparator to generate a full-swing digital signal in an ADC where $V_{out} = V_s$. We wish to characterize the latch in terms of its power consumption and error-rate.

Error-free latch operation occurs when the sampled voltage at the end of the sampling phase, i.e., the *seed voltage*, is large enough to develop into a full-rail swing during the regenerate phase. The behavior of the latch is determined by: (a) input signal swing V_{sw}, (b) sampler bandwidth B, (c) open-loop gain of the regenerator, and (d) the data rate. Latch errors can be classified as: (a) *metastability errors*, and (b) *memory errors*. Metastability errors occur when the seed voltage is not large enough to regenerate into a full-swing signal at the end of the regenerate phase even though the seed voltage has the correct polarity. The metastability issue is well-known [14], and can be compensated for by employing a cascade of regenerators. Memory errors occur when the input signal swing V_{sw} or the sampler bandwidth B is insufficient to obtain a seed voltage with the correct polarity. Unlike metastability errors, memory errors cannot be corrected by cascading additional regenerators.

The transient behavior the latch in Fig. 6(b) is plotted in Fig. 7(a) in a $1.2V$, $65nm$ CMOS process assuming a data rate of $10Gb/s$. An input sampler with $B = 5.3GHz$ was considered, with input swings $V_{sw} = 600mV_{ppd}$ and $V_{sw} = 800mV_{ppd}$.

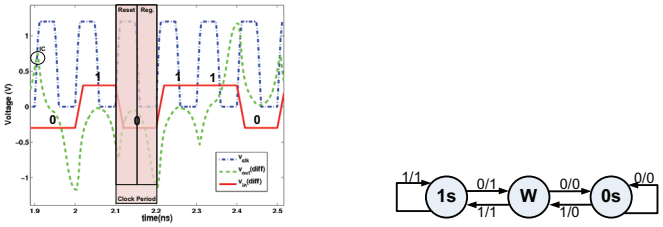

(a) (b)

Fig. 7. Latch modeling: (a) transient waveforms with initial $V_{out} = 0.75V$, an input sequence of 010110, and an input swing $V_{sw} = 600mV_{ppd}$, and (b) the Markov model.

For $B \geq 6GHz$, the latch is memory error-free, i.e., the seed voltage has correct polarity. For lower values of B, memory errors begin to appear as the sampler is unable to track the input signal.

Figure 7(a) depicts the signal transients in and highlights the reset and regenerate phases. We observe that it takes at least two consecutive identical digits (CIDs) for the latch to transition from a strong '1' to a '0', and vice-versa. When the input bits alternate, the output retains polarity. Three latch errors occur, as shown in Fig. 7(a). A Markov model can be employed to capture the behavior of a latch exhibiting memory errors. It turns out that this model needs to be V_{sw}-dependent. The Markov chain transitions corresponding to the input swing of $V_{sw} = 600mV_{ppd}$ is shown in Fig. 7(b). The model is verified by comparing the latch output from circuit simulations with the predicted output bits obtained from the Markov model. For a 1000 bit long PRBS sequence which was verified to have every n-bit pattern up to $n = 10$, the model in Fig. 7(b) predicts all 326 errors perfectly in both number and position of errors. The PRBS sequence also guarantees that all the states and transitions are realized frequently.

B. CAD Challenges

The design of communications-inspired systems (both SAMS and stochastic platforms) presents a number of CAD challenges in all aspects of the design methodology. Some of these are:

1) **Engineering Error-Statistics**: Develop macro-level design and synthesis techniques for digital circuits that result in: a) a graceful increase in error-rates, b) a desired error PMF or error rate, c) and uncorrelated and/or independent errors.
2) **Variability and Noise-aware Modeling**: Automatic macromodeling that captures the impact of variability and noise in device and lithography parameters, and propagating these to the circuit, architecture and system level. Models need to capture behavior outside the traditional envelopes.
3) **Verification and Test**: Define appropriate verification and test metrics for SAMS and stochastic platforms designed using unreliable components.
4) **Compilation and Synthesis**: Develop techniques to aid design space exploration, and automatic synthesis of SAMS-based systems. Develop techniques to efficiently map applications on to programmable stochastic processors.

V. Conclusion

As semiconductor technology marches into the deep nanoscale regime, it is inevitable that stochasticity of the device and circuit fabric will need to be grappled with by both the design and CAD communities. This paper has described SAMS as an elegant design philosophy for the design of reliable and energy-efficient AMS circuits and systems. Hampering the design and deployment of such systems is a lack of appropriate CAD support. We hope that the CAD community will take up this challenge so that the semiconductor industry can continue to reap the benefits of Moore's law for many years to come.

Acknowledgment

This work was supported the Gigascale Systems Research Center (GSRC), one of the six centers funded under DARPA and SRC's Focus Center Research Program (FCRP), and TxACE Task 1836.019. The authors gratefully acknowledge discussions with their colleagues Professor Elyse Rosenbaum, Minwei Lu, Rajan Narasimha, and Andrew Bean.

References

[1] Intl. technology roadmap for semiconductors 2008 update.
[2] O. Agazzi and et al. A 90nm cmos dsp mlsd transceiver with integrated afe for electronic dispersion compensation of multi-mode optical fibers at 10gb/s. In *International Solid-State Circuits Conference*, pages 232–233, Feb. 2008.
[3] H.-M. Bae, J. B. Ashbrook, J. Park, N. R. Shanbhag, A. C. Singer, and S. Chopra. MLSE receiver for electronic dispersion compensation of OC-192 fiber links. *IEEE Journal of Solid State Circuits*, 41(11):2541–2554, 2006.
[4] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarazi, and V. De. Parameter variations and impact on circuits and microarchitecture. In *Proc. of DAC.*, pages 338–342, June 2003.
[5] M. Harwood and et al. A 12.5 Gb/s SerDes in 65nm CMOS using a baud-rate ADC with digital RX equalization and clock recovery. In *IEEE International Solid-State Circuits Conference*, 2007.
[6] R. Hegde and N. R. Shanbhag. Soft digital signal processing. 9(6):813–823, Dec. 2001.
[7] M. Lu, N. Shanbhag, and A. Singer. Ber-optimal adcs for communication links. In *IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1029–1032, June 2010.
[8] B. Murmann. Trends in low-power digitally assisted a/d conversion.
[9] R. Narasimha and N. R. Shanbhag. Design of energy-efficient high-speed links via forward error-correction. 57(5):359–363, May 2010.
[10] P. Schvan and et al. A 24GS/s 6b ADC in 90nm CMOS. In *IEEE International Solid-State Circuits Conference*, 2008.
[11] N. Shanbhag, R. Abdallah, R. Kumar, and D. Jones. Stochastic computation. In *47th DAC, Anaheim, June 2010*, pages 859–854, June 2010.
[12] A. Singer, A. Bean, and J. Choi. Mutual information and time-interleaved analog-to-digital conversion. In *Proc. IEEE Workshop on Information Theory and its Applications*, pages 1–5, Jan. 2010.
[13] G. V. Varatkar, S. Narayanan, N. R. Shanbhag, and D. Jones. Sensor network-on-chip. In *2007 Intl. Symp. on System-on-Chip*, pages 1–4, Nov. 2007.
[14] H. J. M. Veendrick. The behavior of flip-flops used as synchronizers and prediction of their failure rate. *IEEE Journal of Solid State Circuits*, SC-15(2):169–176, 1980.