

ENABLING SYSTEM-LEVEL PLATFORM RESILIENCE THROUGH EMBEDDED DATA-DRIVEN INFERENCE CAPABILITIES IN ELECTRONIC DEVICES

Naveen Verma¹, Kyong Ho Lee¹, Kuk Jin Jang¹ and Ali Shoeb²

¹Department of Electrical Engineering, Princeton University, Princeton, NJ 08544

²Massachusetts General Hospital (now with WeatherBill, San Francisco, CA 94107)

{nverma¹, kyonglee¹, kjang¹}@princeton.edu, ashoeb@alum.mit.edu²

ABSTRACT

Advanced devices for embedded and ambient applications represent one of the most compelling classes of electronic systems, but they also impose more severe constraints on system resources than ever before. Although platform non-idealities have always posed a fundamental limitation, the overheads of conventional margining are now reaching intolerable levels. We describe an alternate approach to hardware resilience that applies to applications where advanced modeling and inference capabilities are required, a rapidly increasing emphasis in many applications. We show how a data-driven modeling framework for analyzing application data can also be used to effectively model and overcome a broad range of hardware non-idealities. Specific examples for biomedical sensors are shown that are able to retain performance with minimal on-line overhead despite the presence of severe digital- and analog-circuit non-idealities.

Index Terms- Hardware resilience, machine learning, stochastic computation, digitally-assisted analog, biomedical devices.

I. INTRODUCTION

Non-ideal behavior of hardware has always been one of the fundamental limitations in electronic systems. The sources of such behavior have been extremely diverse, ranging from imperfect control of transistor processing, to explicit faults in manufacturing, to dependencies on aging and environmental factors. A key challenge is that the circuit- and system-level manifestations have been equally diverse. Although methods have been adopted for handling these, they have relied heavily on various forms of margining that have, in turn, meant sacrificing critical system-level resources. For example, precision analog instrumentation has relied on stabilizing closed-loop topologies with high gain, leading to high power consumption [1]; high- and medium-resolution data conversion has required precise matching of components, increasing both area and power [2]; variation, affecting the noise margins of digital logic and memory, has required setting the supply and threshold voltages at values considerably off the minimum-energy point [3]; and defects in nano-scale lithography have limited the scaling of technology features [4].

While margining has traditionally been viable in electronic systems, some of the most compelling emerging applications face more severe constrictions on system resources (energy, size, etc.) than ever before *and* simultaneously exacerbate the sources of hardware non-idealities. The overhead of conventional margining is thus reaching intolerable levels [5]. The search for alternate approaches to hardware resilience has consequently become one of the most important focus areas, mandating technology-level solutions in the ITRS [6] as well as all methods of circuit-level solutions.

In this paper, we present an approach for overcoming hardware non-idealities using data-driven modeling methods in

applications that require classification of data. Data-driven hardware resilience (DDHR) has the potential to handle severe manifestations from a wide possibility of sources, ranging from imperfect device and component characteristics to explicit manufacturing faults. The approach is based on the observation that, with the increasing importance of recognition, mining, and synthesis applications [7], there is a growing need to enable advanced inference capabilities in electronic systems. The idea is that, in addition to analyzing application data, these capabilities can be exploited to model and overcome the manifestations of static non-idealities that affecting the hardware. Although the manifestations can be severe and nearly impossible to predict a priori, we demonstrate that they can potentially be overcome via data-driven modeling through a supervised machine-learning framework. We illustrate the approach on four types of non-idealities: (1) bit errors in a memory, (2) switching (stuck-at) faults in digital logic, (3) gain non-linearity in an instrumentation amplifier, and (4) integral non-linearity (INL) in an ADC.

II. BACKGROUND

Approaches have emerged that exploit system- and architecture-level frameworks for overcoming non-idealities potentially much more efficiently than conventional margining. Several strong examples exist of analog circuits that take advantage of digital assists in mixed-signal architectures [8]. For example, modeling methods, though applied to a specific non-ideality, have already demonstrated the possibility of substantial linearity improvement in an ADC [9]. As another example, closed-loop calibration of the discrete values possible in a digital modulator have allowed RF transmitters to overcome power-amplifier non-linearities [10].

For digital circuits, approaches have emerged that avoid severe margining overheads by permitting errors and then employing error detection and correction through enhancements at the micro-architecture level of a processor pipeline [11]. Alternatively, approaches, termed *stochastic computation* [12] mitigate hardware overheads by exploiting error tolerance according to application-level metrics (e.g., using differing error characteristics in redundant estimators or soft-voters to reduce the likelihood of severe output errors). Programmable architectures have also been demonstrated that explicitly retain error-free control flow through reliable cores while employing relaxed-reliability cores for computation [13].

The approach presented next leverages the powerful data-driven modeling capabilities enabled by machine-learning methods as a generic way to overcome the manifestations of digital and analog non-idealities. The demonstrations suggest that errors with potentially large and highly irregular distributions can be handled effectively.

III. OVERVIEW OF DATA-DRIVEN HARDWARE RESILIENCE (DDHR)

Data-driven modeling implies the use of data from the application

to construct a model. From this, a decision function can then be derived to perform inference. The approach of DDHR involves forming the model using data that has been processed by the platform and that is thus explicitly affected by the non-idealities of concern. Consequently, DDHR depends strongly on the effectiveness with which the manifestations, which may be physically complex, can be modeled through training. This implies that the error sources must be static. A key to the approach is the use of powerful modeling capabilities offered by supervised machine-learning methods. It is worth noting, that these methods have begun to be incorporated in energy-constrained devices for advanced signal analysis [14, 15].

Fig. 1 illustrates the DDHR concept using a support vector machine (SVM) for binary classification in a sensing device. SVMs are a popular and versatile machine-learning framework for modeling and classification [16]. They analyze data segments by deriving features and representing these as a vector. Feature vectors that have been annotated with class labels form a training set, and instances, called the support vectors, are selected to model the data and compute a decision function for real-time classification. A typical decision function, called the radial basis function (RBF) is shown in Eq. 1 (\vec{x} represents a feature vector of the input signal, \vec{sv}_i represent the support vectors, $K(\cdot)$ is a transformation function, and σ, α_i, y_i, b are training parameters):

$$(\text{Classification}) = \text{sgn} \left(\sum_{i=1}^N K(\vec{sv}_i, \vec{x}) \alpha_i y_i - b \right) = \text{sgn} \left(\sum_{i=1}^N \exp \left(-\frac{\|\vec{sv}_i - \vec{x}\|^2}{\sigma^2} \right) \alpha_i y_i - b \right) \quad (1)$$

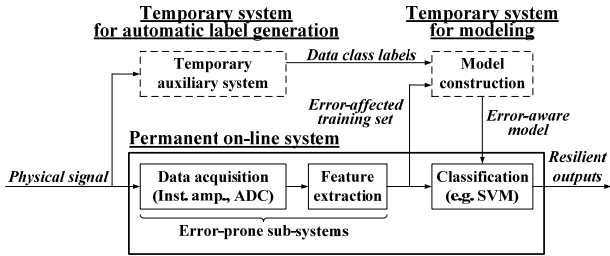


Fig. 1: Illustration of DDHR concept. Temporary systems for modeling and labeling generate an error-aware model from error-affected training data. The error-aware model is then used by a classifier.

With DDHR, the idea is to *derive a training set from feature vectors that are obtained via the non-ideal hardware*. This results in an *error-aware model*. The model is then applied using an error-free classification kernel; while ensuring error-free classification incurs the conventional overhead, this can be offset by employing specialized classifier optimizations [17, 18], thus allowing non-idealities in all preceding stages to potentially be handled efficiently and generically.

If the error sources are static, suitable error-affected data to construct the model is readily obtained by processing data through the device. In addition, however, proper class labeling of the training-set feature vectors is also required. This can be achieved via an expert (in particular, by using active-learning methods [19] to substantially reduce the labeling effort), or it can be achieved automatically by using a temporary error-free auxiliary system to derive the labels (in fact, a low error-mode may be available within the existing system, for instance, by temporarily increasing the supply voltage in a low-voltage device); this approach has shown nearly identical performance to an error-aware system based on perfect, expert labeling [20].

Fig. 2 shows the concept of the error-aware model in an actual application using real data. Initially, the decision function based on the support-vector model is derived from the variances of the physical signal (Fig. 2a). In the presence of hardware errors,

however, the distribution of the feature vectors is altered (Fig. 2b). Training using an error-affected dataset leads to the error-aware model. Fig. 3 shows histograms of the decision function computation (Eq. 1). While the baseline detector (with no errors) can reliably discriminate between the classes, the errors degrade the separation achievable. With the error-aware model, however, the separation between the outputs is effectively restored.




Fig. 2: EEG-signal feature data for a seizure detector (plotted in 2D using principle component analysis) to show (a) the error-free distribution and (b) the distribution after introducing errors (in 20% of the memory cells).




Fig. 3: Histograms of the decision-function output for a seizure detector (a) without errors (baseline), (b) with memory errors, and (c) with errors, but using an error-aware model.

It is worthwhile to note that the performance of an SVM has been shown to degrade gracefully in the presence of low-level feature-vector noise. For instance, Fig. 4 shows the performance as white noise with increasing PSD is applied to the input signal. With low-level bit errors, some degree of resilience is thus retained even without an error-aware model [21]. In general, however, errors induced by hardware non-idealities cannot be assumed to affect only low-order bits (as will be shown in the distributions of Section IV).




Fig. 4: Performance degradation of an SVM is initially graceful when white Gaussian noise is applied to the input EEG signal of a seizure detector (noise power is normalized to the signal power).

IV. DEMONSTRATIONS

In this section, we demonstrate the DDHR approach through the use of two biomedical monitoring applications that utilize real clinical data. Biomedical detectors are representative of advanced sensing applications in that they perform analysis over physically-complex signals; the application signals thus motivate the need for data-driven methods [22]. The applications considered are (1) epileptic seizure detection based on 18-channel electroencephalogram (EEG) sensing [23] and (2) cardiac arrhythmia detection based on electrocardiogram (ECG) sensing [24]. Data is obtained from the MIT-BIH [25] and CHB-MIT [25, 26] databases, respectively. Block diagrams of the detectors are shown in Fig. 5, illustrating the feature computations involved.




Fig. 5: Block diagrams for (a) seizure detector (one EEG channel shown; up to 18 channels can be used) and (b) arrhythmia detector.

The results below consider digital and analog circuits. The digital faults are (A) bit-cell errors in a memory and (B) switching errors in logic gates (stuck-at faults), both of which represent prominent error sources in low-energy designs [3]. The analog non-idealities are (C) gain non-linearity in an instrumentation amplifier and (D) integral non-linearity (INL) in an ADC, which are common data-acquisition errors. The performance is evaluated via the detector’s true-positive (TP) and true-negative (TN) rates, which must both be high, and the false-positive (FP) and false-negative (FN) rates, which must both be low (though all four metrics are examined, results are provided for only the ones where degradations are observed). For all cases, the baseline classifier model (no error-aware) is constructed from a feature-vector training set derived from the nominal processing of Fig. 5; the error-aware model is constructed from a feature-vector training set derived from the error-affected implementations described below.

A. Memory Bit-cell Errors

The feature computation blocks are implemented in Verilog and synthesized to a gate-level netlist. The filter coefficients are assumed to be stored in memory. To inject bit errors, bits are selected at random and their values are flipped. The TP and TN rates for the seizure and arrhythmia detectors are shown in Fig. 6 with respect to increasing fault injection rate (the resulting bit-error rates in the computed features are also shown in parenthesis). The error-aware model retains performance to very high error rates while the performance of the original model degrades rapidly [20].




Fig. 6: Performance of (a) seizure detector and (b) arrhythmia detector.




Fig. 7: (a) RMS error by feature (normalized to the RMS of the feature value) and (b) representative error distributions from four features (cases shown correspond to 20% memory bit-cell errors in the seizure detector).

For illustration, Fig. 7 shows sample feature error characteristics as a result of the injected faults. The magnitude of the errors (Fig. 7a) is large and their distributions (Fig. 7b) are highly irregular, indicating the high flexibility offered by DDHR.

B. Logic-switching Errors (stuck-at faults)

To inject logic switching faults, the synthesized gate-level netlists are manually altered by randomly selecting nodes in the circuit and assigning them a static value of logic ‘1’ or logic ‘0’ with 50% probability. Fig. 8 shows the performance results over ten runs; with logic-switching faults, the performance depends strongly on the precise nodes that are affected (although multiple runs are performed for the memory faults as well, the performance from run-to-run is more consistent, allowing the results to be consolidated into the profiles shown in Fig. 6). The error-aware model is able to consistently restore performance up to fault rates of 0.01% for the seizure detector (corresponding to over 5 nodes in the circuit) and fault rates of 7% for the arrhythmia detector (corresponding to over 100 nodes in the circuit) [20].

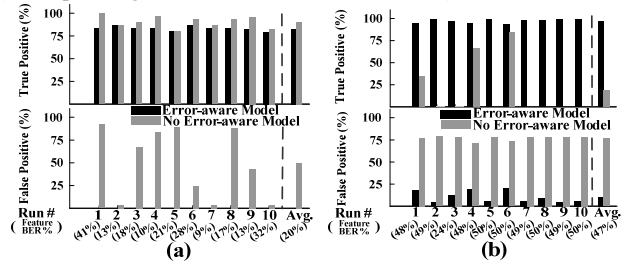


Fig. 8: (a) Seizure detector performance with fault rate of 10^{-4} errors/node (10 runs), and (b) Arrhythmia detector performance with fault rate of 7×10^{-2} per node (10 runs). The error-aware model retains a low FP rate.

C. Instrumentation Amplifier Non-linearity

Instrumentation-amplifier non-linearity is a critical challenging in sensing devices, limiting the dynamic range and ultimately the power consumption of the data-acquisition sub-system. Although a variety of effects on the transfer function are possible, for demonstration, we consider a transfer function given by the following [27]:

$$V_{out} \propto \frac{\alpha V_{in} - \frac{1}{8}(\alpha V_{in})^3 - \frac{1}{128}(\alpha V_{in})^5 - \dots}{\alpha} \quad (2)$$

Here, the output (V_{out}) is set by a linear and several non-linear terms dependent on the input (V_{in}). In this formulation, the ideal output can be normalized to a range of ± 1 , and α can be used to represent a compression factor (with higher levels of α yielding increasingly severe non-linearity). The transfer function that results from this is shown in Fig. 9.




Fig. 9: Transfer curve resulting from the amplifier non-linearity considered.

To analyze amplifier non-linearity, the data is processed through the non-linear transfer function (using terms up to the 5th order). Fig. 10 shows the resulting performance of the seizure detector with respect to α . As with the digital faults, the error-aware model is able to retain the performance to very high levels of the non-ideality. It is worthwhile to note that since analog non-idealities have a gradual effect, the degradation in the classifier performance without the error-aware model is graceful, according to the inherent resilience of an SVM illustrated in Fig. 4 for low-level errors (i.e., this can be seen for $\alpha \leq 100$).




Fig. 10: Seizure detector performance with amplifier non-linearity using (a) TP rate to represent sensitivity and (b) TN rate to represent specificity.

D. ADC Integral Non-linearity

Depending on the ADC topology used, INL errors typically occur with a particular characteristic originating from non-idealities in the circuit components [28]. For demonstration, we consider the characteristic shown in Fig. 11, where INL steps, occurring at the MSB, MSB/2, and MSB/4 transitions, are parameterized by s (this characteristic is common for successive-approximation and pipeline ADCs [28, 29]).




Fig. 11: INL characteristic considered (parameterized by the MSB step, s).

The INL error characteristic with varying step sizes, s , is converted into an ADC transfer curve and then applied to the recorded physiological data in the application. Fig. 12 shows the resulting performance of the seizure detector. As with the previous cases, the error-aware model is able to restore performance even with high levels of the non-ideality.




Fig. 12: Seizure detection performance with ADC INL errors using (a) TP rate to represent sensitivity and (b) TN rate to represent specificity.

V. CONCLUSION

Despite the complex and potentially severe consequences of hardware non-idealities in electronic devices, robust and flexible data-driven methods based on machine learning can overcome non-idealities by effectively modeling the error manifestations. Data-driven training using data affected by the non-idealities leads to an error-aware model that allows the detectors to retain overall performance despite severe errors. Signal classification is a key need in many emerging embedded applications, and it can directly take advantage of an error-aware model with no on-line overhead. We demonstrate this in biomedical detectors that analyze physiological signals using a support-vector machine (SVM) classifier. The advanced modeling capabilities of the SVM framework are shown to provide a highly effective and generic methodology for overcoming potentially severe errors due to digital and analog hardware non-idealities.

ACKNOWLEDGMENTS

We gratefully thank Prof. N. Shanbhag for valuable input. Funding is provided by the Gigascale Systems Research Center, one of six research centers funded under the Focus Center Research Program

(FCRP), a Semiconductor Research Corporation entity.

REFERENCES

- [1] D. Johns and K. Martin, *Analog integrated circuit design*. Hoboken, NJ: Wiley, 1996.
- [2] P. Kinget, "Device mismatch and tradeoffs in the design of analog circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 6, pp. 1212–1224, Jun. 2005.
- [3] N. Verma, et al., "Nanometer MOSFET variation in minimum energy subthreshold circuits," *IEEE Trans. Electron Devices*, vol.55, no.1, pp.163-174, Jan. 2008
- [4] J. Hartmann, "Towards a new nanoelectronic cosmology," *ISSCC*, Feb. 2007.
- [5] M. Horowitz, et al., "Scaling, power, and the future of CMOS," *IEDM*, Dec. 2005.
- [6] International technology roadmap for semiconductor 2010 update. ITRS. <http://www.itrs.net/Links/2010ITRS/Home2010.htm>
- [7] P. Dubey, "Recognition, mining and synthesis moves computers to the era of tera," *Technology at Intel Magazine*, Feb. 2005
- [8] B. Boris Murmann, "A/D converter trends: Power dissipation, scaling and digitally assisted architectures," in *IEEE Custom Integrated Circuits Conference*, Sept. 2008, pp. 105–112.
- [9] B. Murmann and B. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2040–2050, Dec. 2003.
- [10] J. L. Dawson and T. H. Lee, "Cartesian feedback for RF power amplifier linearization," in *Proc. Amer. Control Conf.*, 2004, vol. 1, pp. 361–366
- [11] D. Ernst, et al., "Razor: A low-power pipeline based on circuit-level timing speculation," *Int. Symp. on Microarchitecture*, Dec. 2003.
- [12] N. Shanbhag, et al., "Stochastic computation, DAC, June 2010.
- [13] L. Leem, et al., "Error-resilient system architecture for probabilistic applications," *IEEE/ACM DATE*, Mar. 2010.
- [14] N. Verma, et al., "A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system," *IEEE J. Solid-State Circuits*, vol.45, no.4, pp.804-816, April 2010.
- [15] A. Shoeb, et al., "A micropower support vector machine based seizure detection architecture for embedded medical devices," in *Proc. IEEE Eng. Med. Biol. Soc. Conf.*, Sep. 2009, pp. 4202–4205.
- [16] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel Based Learning Methods*. Cambridge, UK: Cambridge University Press, 2000.
- [17] K. H. Lee, et al., "Improving kernel-energy trade-offs for machine learning in implantable and wearable biomedical applications," *ICASSP*, May 2011.
- [18] K. M. Khan, et al., "Hardware-based support vector machine classification in logarithmic number systems," *ISCAS*, May 2005.
- [19] K. Brinker, "Incorporating diversity in active learning with support vector machines," *Proc. of Int. Conf. on Machine Learning*, Aug. 2003.
- [20] K. H. Lee, et al., "A data-driven modeling approach to stochastic computation for low-energy biomedical devices," *EMBC*, Aug. 2011.
- [21] V. K. Chippa, et al., "Scalable effort hardware design: exploiting algorithmic resilience for energy efficiency," *DAC*, June 2010.
- [22] N. Verma, et al., "Data-driven approach for computation in intelligent biomedical devices: a case study of EEG monitoring for chronic seizure detection," *Journal of Low Power Electronics and Applications*, Vol. 1, Issue 1, June 2011
- [23] A. Shoeb and J. Guttag, "Application of machine learning to epileptic seizure detection," *Proc. of Int. Conf. on Machine Learning*, June 2010.
- [24] E. Ubeyli, "ECG beats classification using multiclass support vector machines with error correcting output codes," *Digital Signal Processing*, May 2007.
- [25] PhysioNet. <http://www.physionet.org>.
- [26] A. L. Goldberger, et al., PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, 101(23):e215-e220.
- [27] P. Wambacq and W. M. Sansen, *Distortion Analysis of Analog Integrated Circuits*. Boston, MA: Kluwer, 1998.
- [28] K. Gulati, "A low-power reconfigurable analog-to-digital converter," Ph.D. dissertation, MIT, MA, 2001.
- [29] N. Verma, "An ultra low power ADC for wireless micro-sensor applications," M.S. thesis, MIT, MA, 2005.