33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

A Data-driven Modeling Approach to Stochastic Computation for Low-energy Biomedical Devices

Kyong Ho Lee, Kuk Jin Jang, Ali Shoeb, and Naveen Verma

Abstract—**Low-power devices that can detect clinically relevant correlations in physiologically-complex patient signals can enable systems capable of closed-loop response (e.g., controlled actuation of therapeutic stimulators, continuous recording of disease states, etc.). In ultra-low-power platforms, however, hardware error sources are becoming increasingly limiting. In this paper, we present how data-driven methods, which allow us to accurately model physiological signals, also allow us to effectively model and overcome prominent hardware error sources with nearly no additional overhead. Two applications, EEG-based seizure detection and ECG-based arrhythmia-beat classification, are synthesized to a logic-gate implementation, and two prominent error sources are introduced: (1) SRAM bit-cell errors and (2) logic-gate switching errors ('stuck-at' faults). Using patient data from the CHB-MIT and MIT-BIH databases, performance similar to error-free hardware is achieved even for very high fault rates (up to 0.5 for SRAMs and $7x10^{-2}$ for logic) that cause computational bit error rates as high as 50%.**

I. Introduction

Recently, biomedical sensors and stimulators have emerged, offering unprecedented modalities for both acquiring physiological signals and for chronically delivering therapy (e.g., conformal implanted/surface electrode arrays, deep-brain stimulators, etc.). In order to take advantage of these in intelligent, closed-loop systems, computational platforms are required that allow high order models of the physiological signals to be applied so that specific states of interest may be detected. The challenge is that these platforms must operate robustly and at highly constrained power levels (i.e., 10-100μW for implantable devices and 1-10mW for wearable devices).

In ultra-low-power platforms, however, hardware errors are becoming increasingly prominent [1]. These originate from transistor variability, manufacturing limitations, device aging, transient events, etc. For instance, aggressive supply-voltage reduction minimizes energy in digital circuits but rapidly increases SRAM, logic-switching, and timing errors [1]. Conventional margining approaches against these leads to intolerable overheads [2], and thus alternate methods have recently emerged for low-power sensors to more efficiently handle hardware errors. These have collectively been referred to as *stochastic computation* [3]. They, for instance, have involved micro-architecture enhancements in a processor to include stages in the pipeline for detecting errors and then re-performing the computation when these occur [4]. Other approaches use redundancy, employing either estimators or soft voters with differing error characteristics in order to reduce the likelihood of errors [3]. Programmable platforms have also emerged that notice the need to retain reliability in the computation control flow; these incorporate reliable cores for error-free control while permitting relaxed reliability for the data computation cores [5].

In this paper, we present a data-driven approach to stochastic computation that permits very large error rates to be handled with nearly no additional on-line overhead or redundancy. Data-driven methods have emerged as a powerful approach for modeling patient signals [6]. This has come about first due to the large-scale availability for physiological data in the healthcare domain, and second due to the advancement of machine-learning techniques that provide extremely efficient methods for analyzing large datasets and extracting correlations into accurate models. These methods have recently begun to be incorporated into ultra-low-power devices [7]. We describe how the modeling capabilities of these methods can be exploited to accurately model not only the physiological signals, *but also the specific manifestations of hardware computation errors that affect their processing*. This leads to an *error-aware model* that provides high resiliency even when the fault sources affect high-order bits. We synthesize two applications to a logic-gate level implementation: (1) seizure detection based on electroencephalographs (EEGs); and (2) cardiac-arrhythmia beat classification based on electrocardiographs (ECGs). We introduce hardware errors corresponding to two prominent fault sources, and using patient data from the CHB-MIT [8, 9] and MIT-BIH [8] databases, we demonstrate how detection performance is restored despite high error rates.

II. Data Driven Stochastic Computation

A. Overview

Data-driven stochastic computation is based on the use of machine-learning classifiers for biomedical detection. These (1) model the physiological correlations in the signal biomarkers, and (2) allow that model to be adjusted to account for perturbations in the correlations that result due to hardware computation errors. Fig. 1 shows the structure of data-driven algorithms, illustrated by the representative applications considered in this work. Biomarkers are specific

Manuscript received March 26, 2011. This work was supported in part by the Giga-scale Systems Research Center, one of six research centers funded under the Focus Center Research Program (FCRP), a Semiconductor Research Corporation entity.

K. H. Lee, K. J. Jang, and N. Verma are with Princeton University, Princeton, NJ 08540 USA (phone: 609-258-1424; fax: 609-258-3745; email: kyonglee@princeton.edu).

A. Shoeb is with Massachusetts General Hospital and MIT, Boston, MA 02139, USA (email: ashoeb@mit.edu).

parameters in the patient signals that have some correlation with the physiological states we are interested in detecting. For seizure detection (Fig. 1a), information is contained in the spectral energy distribution of a patient's EEG. Biomarkers thus consist of the spectral energy extracted using a bank of modulated filters (and energy accumulators) applied to each EEG channel [10]. For the arrhythmia detection (Fig. 1b) algorithm, biomarkers consist of a wavelet transform applied to the patient's ECG during each heartbeat segment [11]. In both cases, the biomarkers are then interpreted by treating them as a feature vector for classification via a support vector machine (SVM).

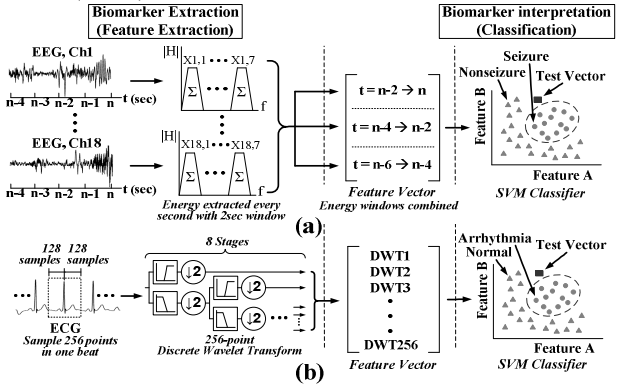

Fig. 1: Algorithms for (a) seizure detection [10] and (b) cardiac-arrhythmia beat classification [11]. The EEG and ECG signal biomarkers form the respective feature vectors for classification.

SVMs are widely used machine-learning classifiers that are popular due to their computational efficiency. They require an initial training phase where pre-labeled feature-vector data is provided and used to derive a set of *support vectors*. The support vectors ($\vec{sv}_i$) represent the decision boundary (illustrated in Fig. 1) through the following computation, where the incoming data (represented by the feature vector $\vec{x}$) is classified based on the resulting sign (a radial basis function (RBF) SVM kernel is shown):

$$Let\ \vec{sv}_i = [sv_{i1}\ sv_{i2} \cdots sv_{iM}]\ and\ \vec{x} = [x_1\ x_2 \cdots x_M]$$

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

(N is the number of support vectors, M is the feature-vector dimensionality, and σ is a training parameter (α_i, y_i, and b are other kernel parameters)).

SVMs allow complex decision boundary modeling over a multi-dimensional information space (i.e., feature space), and they provide efficient methods for developing new models based on labeled datasets. These factors have enabled biomedical detectors with high specificity despite the patient-to-patient variability and numerous physiologic variances exhibited by practical signals. As described below, however, these factors can also be exploited towards modeling and overcoming hardware errors.

B. Modeling Hardware Errors

Fig. 2 illustrates the concept of modeling hardware errors using an SVM. The biomarker features result in clusters in the feature space (e.g., corresponding to seizure and non-seizure data). However, hardware errors in the biomarker extraction computation alter the clusters, as shown in Fig. 2b (for illustration, a 2-D feature space has been derived using principle component analysis (PCA), and errors have been introduced as described in Section III.B). By training the SVM on the new clusters, the classifier performance can be restored, leading to the *error-aware decision boundary*. Since SVMs can offer a high degree of flexibility in the decision boundary over all of the feature dimensions, diverse and large-magnitude changes can be handled. Thus, this approach is not limited by the severity of the errors, but rather by how they degrade the cluster separation between the classes. Section III.C shows that errors of a high rate and magnitude can be overcome for actual patient data.

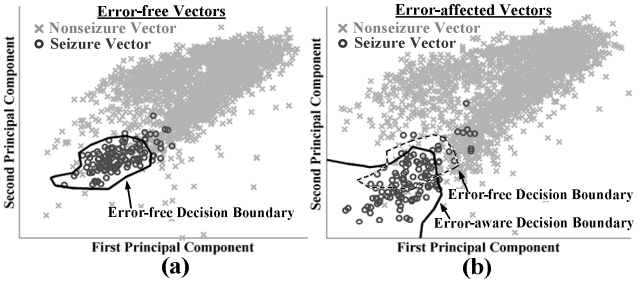

Fig. 2: Scatter plots from PCA of EEG features for seizure detection for the cases (a) without hardware errors and (b) with hardware errors (20% SRAM errors). The SVM decision boundaries are shown, illustrating how an error-aware boundary can restore detector performance.

C. Generating and Applying Error-aware Models

The generation and application of an error-aware model requires (1) error-free classifier kernel computation (i.e., Eq. 1) and (2) training of the classifier using data affected by the actual feature computation errors. The classifier kernel has the benefit that it is standard across applications; feature computation, on the other hand, is variable from application-to-application and, potentially, patient-to-patient due to the range of clinical considerations affecting the precise choice of biomarkers. Low-power implementations of the classifier are thus possible, especially using dedicated hardware [12].

In order to train the classifier, a set of feature vectors is required that have been derived using the error-prone hardware and that have been assigned class labels appropriately. Labeling of the error-affected data can be accomplished by a variety of methods. Conventionally, detectors have relied on patient admission for parameter tuning. However, in Section III.D, we also consider the system in Fig. 3, which uses a temporary auxiliary device for automatic generation of the labels. The auxiliary device does not provide alarms in response to hardware errors, but rather labels based on the error-free processing. Such a device can be realized within the permanent device by temporarily placing it in a low-error (but higher energy) mode (e.g., through the use of dynamic voltage scaling [13]). Then, detector training can proceed on an infrequent or one-time basis.

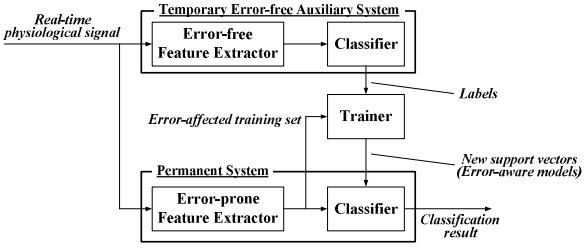

Fig. 3: Error-prone permanent system with temporary error-free auxiliary system to provide labels for the error-aware model generation.

III. Experimental Demonstration

A. Logic-gate Implementation of Feature Processing

The feature-extraction processing for the seizure- and arrhythmia-detection algorithms is synthesized to a library of logic gates for a 130nm IBM CMOS technology. Fig. 4a shows the implementation for the seizure detector. It consists of a decimation filter to down-sample the EEG data for the spectral range of interest. This is followed by seven FIR band-pass filters and absolute-value accumulators to represent the spectral energy over 2 seconds in each bin. This results in 7 features per channel, and the processing is applied over 18 channels and over 3 time windows to form a feature vector of 378 dimensions.

Fig. 4b shows the implementation for the arrhythmia detector. A Daubechies 4 wavelet transform is applied over 256 samples of the ECG. The eight-stages are formed using 4th order FIR filters, and they give a total of 256 features.

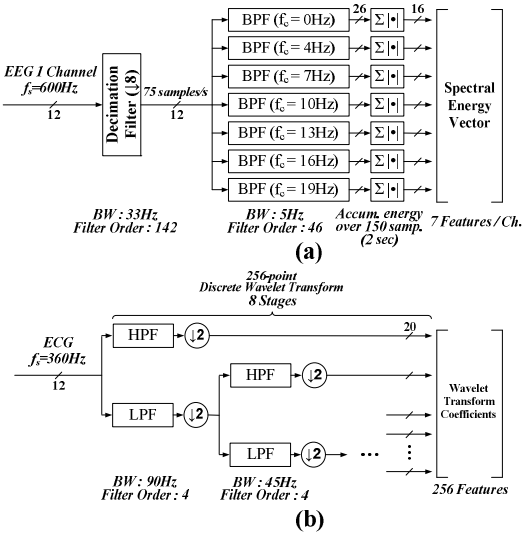

Fig. 4: Feature extraction implementation for (a) seizure detector and (b) arrhythmia detector.

B. Introducing Hardware Errors

Two types of hardware errors that are of primary concern in emerging low-power devices [1] are separately introduced. SRAM errors are introduced by randomly flipping bits in the stored filter-tap values, and logic errors are introduced by randomly inserting stuck-at-1 or stuck-at-0 faults at nodes in the gate-level netlist. These result in high bit-error rates and high error magnitudes. As an example, Fig. 5 shows the RMS error for each feature (normalized to the error-free RMS value of the feature) for the case of 20% SRAM bit-cell errors applied to the seizure detector. For most of the features, the RMS error is larger than the RMS value of the feature itself. Fig. 6 shows the feature error distributions for four representative features, illustrating that, in addition to large magnitude, the errors have highly irregular distributions.

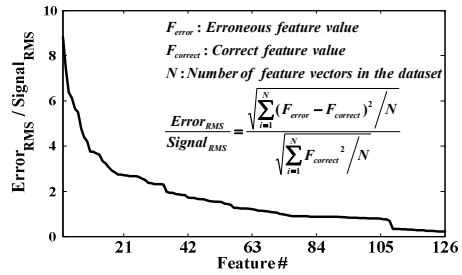

Fig. 5: RMS error by feature (normalized to RMS value of each feature), for 20% SRAM bit-cell errors in the seizure detector. (Note: features are intentionally plotted in order of decreasing normalized error.)

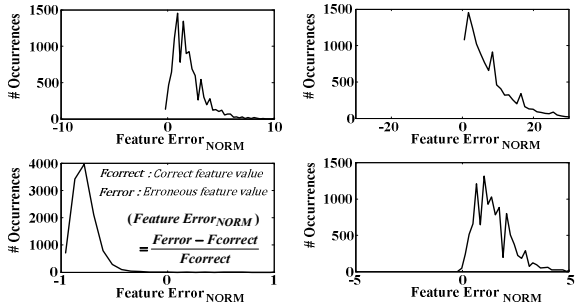

Fig. 6: Representative error distributions for four features for 20% SRAM bit-cell error s in the seizure detector.

C. Performance Results

To evaluate the performance of error-aware modeling, patient data from the CHB-MIT and the MIT-BIH databases are used for the seizure- and arrhythmia-detection applications respectively. The performance depends on the ability of the classifier to discriminate between feature vectors of the positive class (e.g., seizure/arrhythmia data) and those of the negative class. Qualitatively, the benefit of the error-aware model can be seen by the sample histograms of Fig. 7. These correspond to the classifier output for seizure detection in (a) the baseline case (i.e., no errors), (b) the case with errors (10% SRAM bit-cell errors), and (c) the case with errors, but using an error-aware model, showing how the detector's ability to separate the classes is restored.

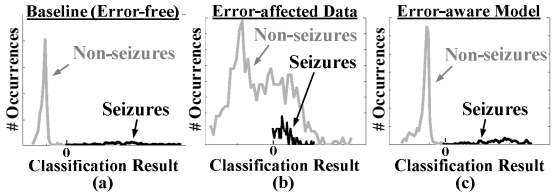

Fig. 7: SVM classification histograms for seizure detection (a) without errors (baseline), (b) with 10% SRAM errors, and (c) with 10% errors, but using an error-aware classifier model.

Fig. 8 shows the performance for seizure detection versus SRAM bit-cell error rate (the resulting BERs in the feature vectors are shown in parenthesis). The sensitivity is represented by the true-positive rate, and the specificity is represented by the true-negative rate. While the sensitivity of

the original detector degrades to nearly zero, that with the error-aware model remains high even at very high error rates.

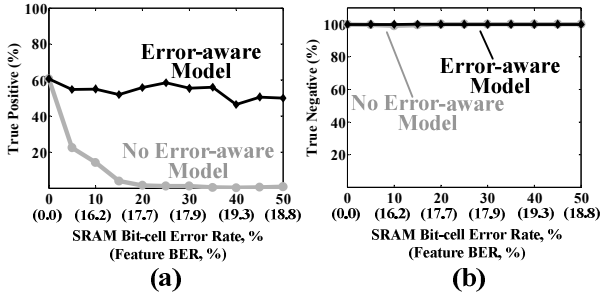

Fig. 8: Seizure detector performance with SRAM bit-cell errors using (a) true-positive rate to represent sensitivity, and (b) true-negative rate to represent specificity. The BER is shown in brackets, though this is not directly correlated with the magnitude of the feature errors nor the ability to recover from errors.

For the logic errors, the error-aware mode was found to consistently restore performance for fault rates up to $1x10^{-4}$ errors/node (corresponding to over 5 nodes in the design). Although higher error rates were sustainable, the actual performance depends on the specific nodes affected, and thus at higher rates, the performance was not consistent from run-to-run. Fig. 9 shows results over 10 runs for the $1x10^{-4}$ case. Here, false-positive rate is used to represent specificity rather than true-negative rate, since this indicates degradations more clearly when the negative class (i.e., non-seizure class) consists of disproportionately more data.

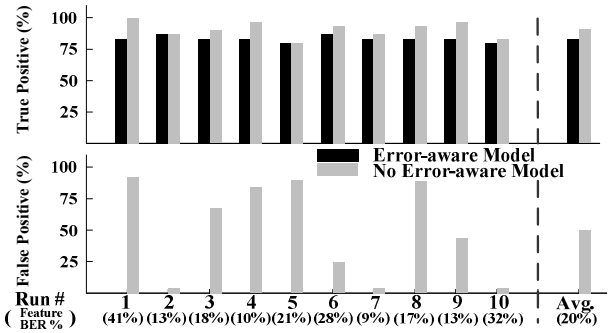

Fig. 9: Seizure detection performance with logic fault rate of $1x10^{-4}$ errors/node (over 10 runs). The error-aware model restores the specificity, as indicated by nearly 0% false positive rate.

Corresponding results for arrhythmia detection are shown in Fig. 10 and 11. Although the sensitivity and specificity do not degrade to zero, both saturate at <50%, indicating no substantial class separation without the error-aware model. For logic faults, rates of $7x10^{-2}$ are consistently overcome.

D. System Demonstration

Instead of pre-labeling to retrain the classifier to an error-aware model, Fig. 12 considers the system in Fig. 3, where the class labels are derived using a temporary, error-free auxiliary system. As shown for both seizure and arrhythmia detection, performance similar to ideal pre-labeling is achieved, and the error-aware model retains performance similar to the error-free case.

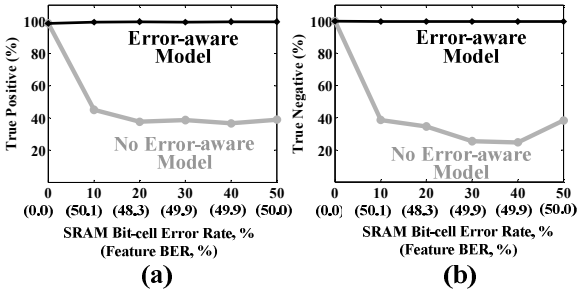

Fig. 10: Arrhythmia detection performance with SRAM bit-cell errors using (a) true positive rate to represent sensitivity and (b) true negative rate to represent specificity.

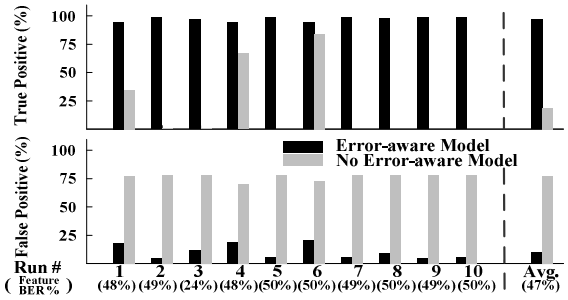

Fig.11: Arrhythmia detection performance with logic error rate of $7x10^{-2}$ per node (10 runs). Error-aware model restores the performance.

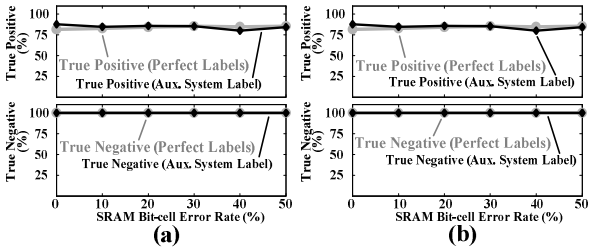

Fig. 12: Comparison of error-aware models generated by perfect labeling with those generated by auxiliary-system labeling for (a) seizure detector and (b) arrhythmia detector having SRAM bit-cell errors.

REFERENCES

[1] 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

[2] M. Horowitz, et al., "Scaling, power, and the future of CMOS," *IEDM*, Dec. 2005.

[3] N. Shanbhag, et al., "Stochastic computation, *DAC*, June 2010.

[4] D. Ernst, et al., "Razor: A low-power pipeline based on circuit-level timing speculation," *Int. Symp. on Microarchitecture*, Dec. 2003.

[5] L. Leem, et al.,"Error-resilient system architecture for probabilistic applications," *IEEE/ACM DATE*, Mar. 2010.

[6] D. Hau, et al., "Learning qualitative models from physiological signals," *AAAI Symp. Artificial Intelligence in Medicine*, 1994.

[7] N. Verma, et al., "A micro-power EEG acquisition SoC with integrated seizure detection processor for continuous patient monitoring," *Symp. VLSI Circuits,* June 2009.

[8] PhysioNet. *http://www.physionet.org.*

[9] A. L. Goldberger, et al., PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, 101(23):e215-e220.

[10] A. Shoeb, et al., "Application of machine learning to epileptic seizure detection," *ICML*, June 2010.

[11] E. Ubeyli, "ECG beats classification using multiclass support vector machines with error correcting output codes," *Digital Signal Processing*, May 2007.

[12] K. M. Khan, et al., "Hardware-based support vector machine classification in logarithmic number systems," *ISCAS*, May 2005.

[13] A. Chandrakasan, et al., "Technologies for ultradynamic voltage scaling" *Proc. of IEEE*, vol. 98, no. 2, pp. 191-214, Feb. 2010.