NURIP: Neural Interface Processor for Brain State Classification and Programmable-Waveform Neurostimulation

Gerard O'Leary, *Student Member, IEEE,* David M. Groppe, Taufik A. Valiante, Naveen Verma, *Member, IEEE,* and Roman Genov, *Senior Member, IEEE*

***Abstract*—The advancement of implantable medical devices for the treatment of neurological disorders demands energy-efficient, low-latency processors for responsive, safe, personalized neuromodulation. A 130nm CMOS neural interface processor (NURIP) is presented to perform brain state classification and closed-loop control using programmable-waveform electrical stimulation. The architecture features an autoencoder neural network for both spatial filtering and dimensionality reduction. Dedicated feature extraction blocks are implemented for univariate (signal-band energy, SE) and multivariate (phase locking value, PLV, and cross-frequency coupling, CFC) neural signal processing. The proceeding exponentially decaying-memory support vector machine (EDM-SVM) accelerator employs these features for hardware-efficient brain state classification with a high temporal resolution. An integrated digitally charge-balanced waveform generator enables flexibility in finding optimal neuromodulation paradigms for pathological symptom suppression. The SoC is validated using the EU human intracranial EEG (iEEG) epilepsy dataset, achieving a seizure sensitivity of 97.7% and a false detection rate of 0.185 per hour while consuming 169 μJ per classification.**

***Index Terms*—Neural Interface Processor, Exponentially Decaying Memory, Support Vector Machine, Phase Locking Value, Cross-Frequency Coupling, Signal Energy, Neuromodulation, Waveform Generation, Binary Exponential Charge Recovery.**

I. Introduction

THE synchronized firing of local neural populations within the brain gives rise to oscillations known as local field potentials, or LFPs. These oscillations are now understood to reflect the underlying activity of the brain, and categorizing their patterns has led to fundamental insights into physiological and pathological neural mechanisms [1]. For example, sleep is a recurring state of mind and body, characterized by altered consciousness, relatively inhibited sensory activity and inhibition of nearly all voluntary muscles. As an individual falls asleep neural activity measured through electroencephalography (EEG) initially transitions from a state of high-frequency, low-voltage waves in the waking state to higher voltage, slower waves representing non-rapid eye movement (NREM) sleep [2]. In the case of pathological conditions such as epilepsy, a seizure can be characterized by abnormal excessive or synchronous EEG activity. The discovery of such anomalous EEG patterns in epilepsy and Parkinson's disease has led to breakthrough therapies based on disrupting the activity associated with seizures and motor tremor respectively using electrical stimulation. However, existing treatments have limited efficacy with only 13% of patients achieving seizure freedom for at least 1 year [3].

Advances in low-power CMOS process technology have revolutionized the capabilities of implantable medical devices. It is now feasible to integrate increasingly more complex and feature-rich digital processing systems with analog circuits for the acquisition of low-amplitude neural signals. Neural signal processing algorithms which extract pathological biomarkers were once only practical to compute on high-power computing systems, but can now be included in a single mixed-signal systems on chip (SoC) [4], [5]. Furthermore, developments in the field of hardware-based machine learning have brought forward efficient algorithm implementations which enable the classification of patient-specific biomarkers for personalized treatments [6] [7]. The combination of these three components: digitization, biomarker extraction and individualized classification enables devices which can measure and identify complex brain dynamics.

Upon the detection of a pathological state, it has been demonstrated that the delivery of an electrical stimulus can be used to influence neural activity and suppress symptoms [8]. However, while existing devices employ a simple bi-phasic pulse waveform, it has been shown that alternatives have the potential to excite networks more selectively and with reduced energy [9] [10]. Furthermore, the brain is a complex dynamical system and neurostimulation devices should be capable of adapting the stimulus in response to changing physiological environments [11]. It is therefore desirable to programmatically synthesize appropriate waveforms in an online manner. However, such waveforms must be constrained to ensure that tissue damage does not arise through the effects of excessive stimulation charge buildup.

As it is now possible to integrate neural signal acquisition, signal processing, machine learning model acceleration and neuromodulation waveform generation on a single SoC, the path is becoming clear for a revolution in the interface between the brain and computational devices. Towards this goal, this paper presents an implantable digital neural interface processor

G. O'Leary and R. Genov are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada. (contact: gerard.oleary@mail.utoronto.ca)

D. M. Groppe is with the Krembil Neuroscience Center, Toronto, ON M5T 2S8, Canada.

T. A. Valiante is with the Department of Surgery (Neurosurgery), University of Toronto, ON M5T 2S8, Canada.

N. Verma is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544.

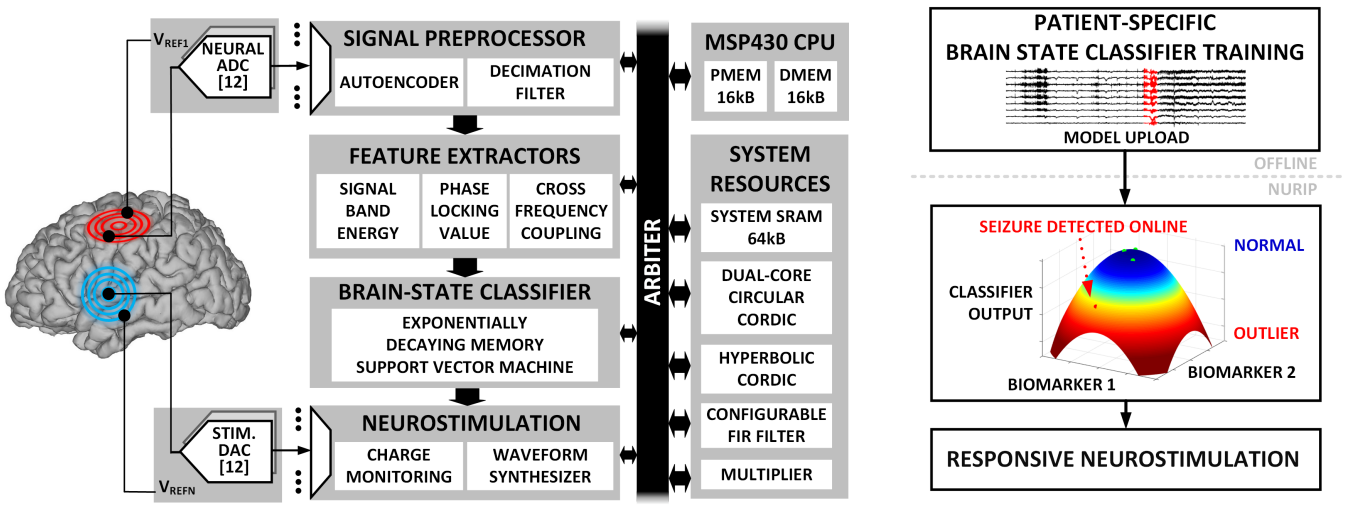

Fig. 1. (a) NURIP system architecture for responsive neuromodulation. The preprocessor stage performs signal conditioning and dimensionality reduction before the extraction of three neural signal biomarkers. The proceeding exponentially decaying memory support vector machine (EDM-SVM) classifies brain states using an offline-trained patient-specific model. Upon the detection of a pathological state, the neuromodulation waveform synthesizer generates an appropriate preventative stimulus. (b) NURIP configuration workflow for responsive neurostimulation based on patient-specific seizure classification.

(NURIP) which performs brain state classification for reliable seizure prediction and contingent seizure abortion. This paper is an extended description of the work introduced in [12], focusing on five key advances:

- The inclusion of an on-chip autoencoder neural network for signal conditioning and dimensionality reduction.
- The integration of a diverse array of univariate and bivariate neural signal processing feature extractors with on-chip machine learning acceleration.
- The implementation of the EDM-SVM for effective and hardware efficient time-series classification.
- On-chip neuromodulation waveform synthesis for precise control of neural activity and online stimulus adaptation.
- The introduction of binary exponential charge recovery (BECR) for digital charge-balanced neurostimulation.

These developments build on existing work from several perspectives. The closed-loop neurostimulation SoC presented in [4] implements the PLV biomarker and manually-set, threshold-based seizure detection. NURIP includes an optimized implementation of this biomarker (Section IV-B) with a 9x reduction in area and a 5x reduction in power. Furthermore, the additional biomarkers used in this work along with data-driven classification improves seizure sensitivity by $> 20\%$.

The general purpose biomedical signal processing platform detailed in [13] accelerates fast Fourier transform (FFT) and signal energy extraction. However, the SoC relies on CPU-based classifier computation, which is orders of magnitude less energy efficient than dedicated hardware accelerators such as the EDM-SVM used in this work [6]. The SVM accelerator introduced in [6] relies on an MSP430 CPU for feature computation, which is up to 754x slower than the dedicated feature extractors presented here (Fig. 6).

The SoC presented in [7] uses SE with a combination of SVM classifiers to tradeoff between detection sensitivity and specificity, but uses a windowing approach with a limited ability to capture complex temporal EEG dynamics (see Section V). The use of SE with PLV, CFC and the EDM-SVM approach in this work results in increased seizure sensitivity and a 45% reduction in false detections (Table V).

The neural-prosthetic device in [14] uses FFT and approximated entropy (ApEn) features along with a linear least squares classifier. However, this device, along with those mentioned previously, relies on bi-phasic pulses for symptom control. NURIP enables the use of on-chip arbitrary waveform generation for precise stimulus control with a digital charge balancing technique to mitigate electrode and tissue damage.

The outline of this paper is as follows. Section II introduces the overall system architecture. Section III describes the preprocessing stages required to condition and manage the recordings from an array of neural recording channels. Section IV details the implemented feature extraction cores including SE, PLV and CFC. Section V outlines the classification of brain states and introduces the exponentially decaying memory support vector machine. Section VI describes the classifier performance with an application demonstration using the EU human iEEG epilepsy database [15]. Section VII outlines the programmable charge-balanced neurostimulation waveform generator. Section VIII describes the VLSI implementation of NURIP and the SoC measurement results. Finally, section IX concludes the paper and describes future directions for implantable brain-state classifiers.

II. System Architecture

The initial preprocessing stage following the digitization of neural signals features an autoencoder neural network for both iEEG spatial filtering and dimensionality reduction. Dedicated feature extraction blocks are implemented for univariate (SE) and multivariate (PLV, CFC) neural signal processing. The

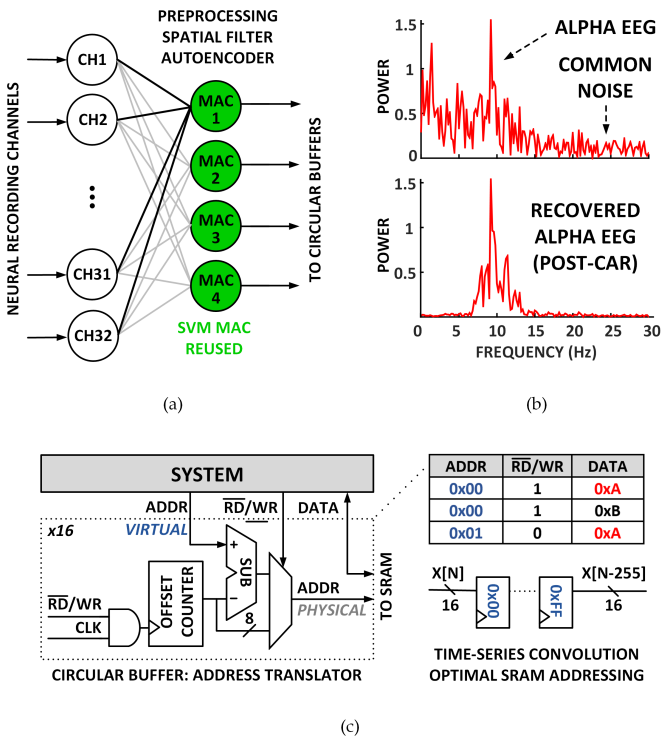

Fig. 2. (a) An autoencoder is implemented with multiplier reuse for input spatial filtering and dimensionality reduction. (b) Experimental results demonstrating the removal of noise from a channel exhibiting alpha-band activity using common average referencing (CAR) [16]. (c) The circular buffer (CB) maintains a continuous window of incoming samples, where each sample is mapped to varying physical SRAM address while the corresponding virtual addresses used by the rest of the system remains fixed.

proceeding EDM-SVM accelerator employs these features for patient-specific brain state classification. A general-purpose CPU facilitates additional custom feature extraction and system control. In response to the detection of a pathological brain state, an appropriate charge-balanced modulation waveform is synthesized to control the operation of an output neurostimulator. The scalable architecture is agnostic to the number and type of analog interface channels to be processed, supporting increasingly high channel counts and new interface paradigms such as optogenetic recording and stimulation.

As the neural signals of interest are typically sampled at low frequencies (256Hz-1KHz for LFPs), the architecture is optimized for reduced power consumption and area over performance. This design consideration can be exploited by time-sharing common system resources such as a configurable-order FIR filter, 32-bit MAC, hyperbolic and circular CORDIC blocks, and system SRAM as shown in Fig. 1.

III. Preprocessing

As the number of integrated recording channels scales to increase the spatial coverage of signal acquisition, there is a corresponding increase in the volume of data which must be processed. Dimensionality reduction can be performed to reduce the required computation while minimizing the loss of information. The preprocessor supports spatial filtering and principal component extraction through the use of an autoencoder neural network. In this case, dimensionality is reduced from 32 recording channels to 4 weighted combinations (such as in principal component analysis), reducing the processing requirements by 8x. The autoencoder is implemented by reusing shared on-chip computing resources which minimizes the area overhead as recording channel arrays scale towards thousands of channels [17].

The autoencoder is an unsupervised learning algorithm that applies backpropagation to train an encoding layer which minimizes the error between an input x_j and the reconstructed output $\hat{x_j}$ as follows:

$$H_i = \sum_{j=1}^{32} W_{ij} x_j + b_i \qquad (1)$$
$$\hat{x_j} = \sum_{i=1}^{4} W_{ij}^T H_i + b_j$$

Where H_i is the encoding hidden-layer node and W_j and b_i are the model parameters. Training is performed offline and the feedforward path can be computed on the implant using the model stored in on-device SRAM. The linear transfer function used is equivalent to principal component analysis (PCA) [18]. This is used to separate multichannel EEG into temporally and spatially independent components that can often be associated with particular neural generators. One drawback of PCA is that it attempts to combine as much of the data variance as possible into each component, even though it may include several temporally independent sources. The biomarkers used in this work (such as PLV, introduced in Section IV) rely on temporal preservation. For this reason, the preprocessor also supports the storage of up to 12 raw channel sample streams.

The autoencoder structure can also be used to implement common average referencing (CAR) for an EEG noise reduction of up to 30% [16]. Fig. 2 demonstrates removal of noise from a channel exhibiting simulated alpha-band activity. In this case, all weights are equivalent and a multiply and accumulate node provides the average of all samples. This average can then be subtracted to reject stimulus artifacts and reduce noise before the feature extraction stage.

The preprocessor is also responsible for data management of incoming sample streams from an array of analog front end ADCs. For time-series analysis, a moving window of the most recently recorded samples must be stored for signal processing purposes. A 256-sample circular buffer (CB) maintains a continuous window of incoming samples (Fig. 2), where each sample is mapped to varying physical SRAM address while the corresponding virtual address remains fixed. This ensures that the latest recording sample is always in a known SRAM location. Preprocessor outputs are stored in 16 CBs, and a further 16 can be enabled for intermediate data management during later signal processing stages. They are mapped to 8-16 kB of address space within 64 kB of global SRAM.

IV. Feature Extraction

There exists a tradeoff in device signal processing between embedded computation and wireless transmission for remote processing. Devices must operate on a low power budget to maximize their battery life [19], and reduce the number of replacement surgeries which result in a risks of infection and an

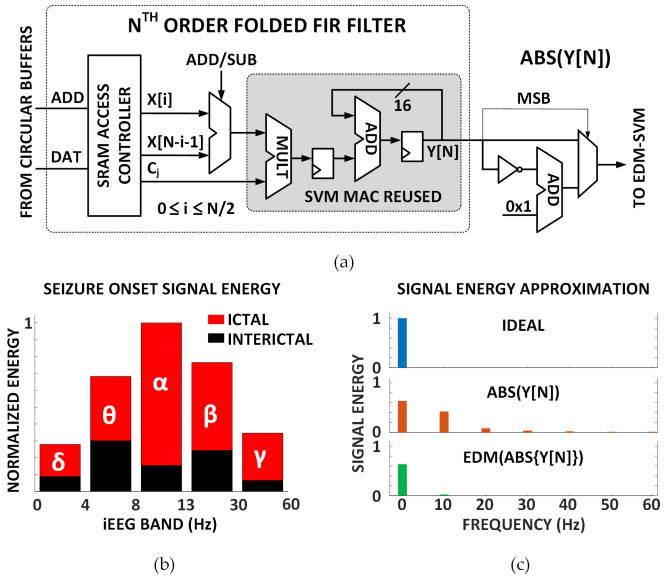

Fig. 3. (a) Signal band energy feature extraction is performed using configurable order Type-I FIR filter. The absolute value a bandpass filter output is taken as a measure of signal energy. (b) The seizure onset energy difference in physiological signal bands for patient FR_1093 in the EU epilepsy database [15]. (c) The error in a 5 Hz signal energy approximation using absolute filter outputs is compensated for in the EDM-SVM classifier (Section V).

additional clinical burden. Furthermore, device designers must consider a tradeoff between thermal limits of heat dissipation and wireless communication [20]. A comparative study found the power consumption of local signal processing an order of magnitude lower than typical wireless data transmission [5].

Following the preprocessor, a subsequent array of three configurable neural signal feature extractors, shown in Fig. 1, enables custom patient-specific processing to maximize classifier performance; Signal Band Energy, Phase Locking Values, and Cross-Frequency Coupling. The ensemble of these three biomarkers yields a high-dimensional feature space for the classifier to distinguish between brain states.

A. Signal Band Energy

The energy in physiological signal bands can be used to characterize brain states based on recorded neural signals. An example is shown in Fig. 3 where elevated energy is seen during seizure (ictal) events across physiological signal bands including δ (0-4 Hz), θ (4-8 Hz), α (8-13 Hz), β (13-30 Hz) and γ (30-60 Hz) when compared to normal (interictal) activity. A configurable order Type-I FIR filtering block reuses system resources to isolate these signal bands. The filter uses shared symmetric coefficients during multiplications to halve the number of operations [22]. The absolute output value of each bandpass filter is taken as a measure of signal energy (Fig. 3). This approximation acts as full-wave rectification and the resulting feature includes noise harmonics at multiples of the extracted band. This is compensated for during the classification stage in Section V.

B. Phase Locking Value

Neural connectivity refers to a pattern of anatomical links between distinct neural populations within the brain. Connectivity patterns are formed by structural links such as synapses or fiber pathways. Neural activity is constrained by connectivity, and quantitative measures are therefore crucial to understanding how the brain processes information [23]. Recent work has led to the discovery of a "preictal state" characterized by a desynchronization of the neuronal populations related to the epileptogenic focus before a seizure onset [8]. Phase locking occurs at specific physiological frequencies which are bandpass filtered for each channel. Synchronization is detected between a channel pair when the difference between the instantaneous phases in the extracted bands, defined as f_0 and f_1, remains constant.

The processing flow required for PLV extraction is shown in Fig. 4 First, a Hilbert FIR filter is applied to bandpass filtered recording sites f_0 and f_1 to obtain an analytic signal. The phase difference between both signals is calculated as:

$$\Delta\phi(t) = \phi_{f_0 i(t)} - \phi_{f_1 i(t)} \tag{2}$$

This angle is used to create an instantaneous complex vector which is constructed using a dual-core COordinate Rotation DIgital Computer (CORDIC) for sine and cosine generation [24]. The magnitude of the average of N vectors is used as a measure of phase locking. If the average $\Delta\phi_i$ is 0, both f_0 and f_1 are phase-locked. FIR-based moving average filters are inefficient in terms of both area and power consumption. Due to the narrow bandwidth of the signal, this moving average can be efficiently replaced by the following IIR approximation:

$$\begin{aligned} Z_{Re}[t] &= \alpha cos(\Delta\phi) + (1-\alpha) Z_{Re}[t-1] \\ Z_{Im}[t] &= \alpha sin(\Delta\phi) + (1-\alpha) Z_{Im}[t-1] \end{aligned} \tag{3}$$

$$\begin{aligned} R &= \frac{1}{N}\sqrt{\left(\sum_{i=0}^{N-1}(cos(\Delta\phi_i)\right)^2 + \left(\sum_{i=0}^{N-1}(sin(\Delta\phi_i)\right)^2} \\ &\approx \sqrt{(Z_{Re}(t))^2 + (Z_{Im}(t))^2} \end{aligned} \tag{4}$$

Two key optimizations have been made to provide power and performance improvements over existing implementations [25]. The IIR approximation presented here results in a 60% decrease in group delay latency. Furthermore, only two CORDIC computations must be performed at a given stage in the processing pipeline. The overall number can thus be reduced from five to two with resource sharing (Fig. 4). These optimizations result in a 5x reduction in power.

C. Cross-Frequency Coupling

A physiological mechanism known as cross-frequency coupling has been identified as playing a key role in biological information processing in the brain [26]. Phase-amplitude coupling is a particular method which has been theorized

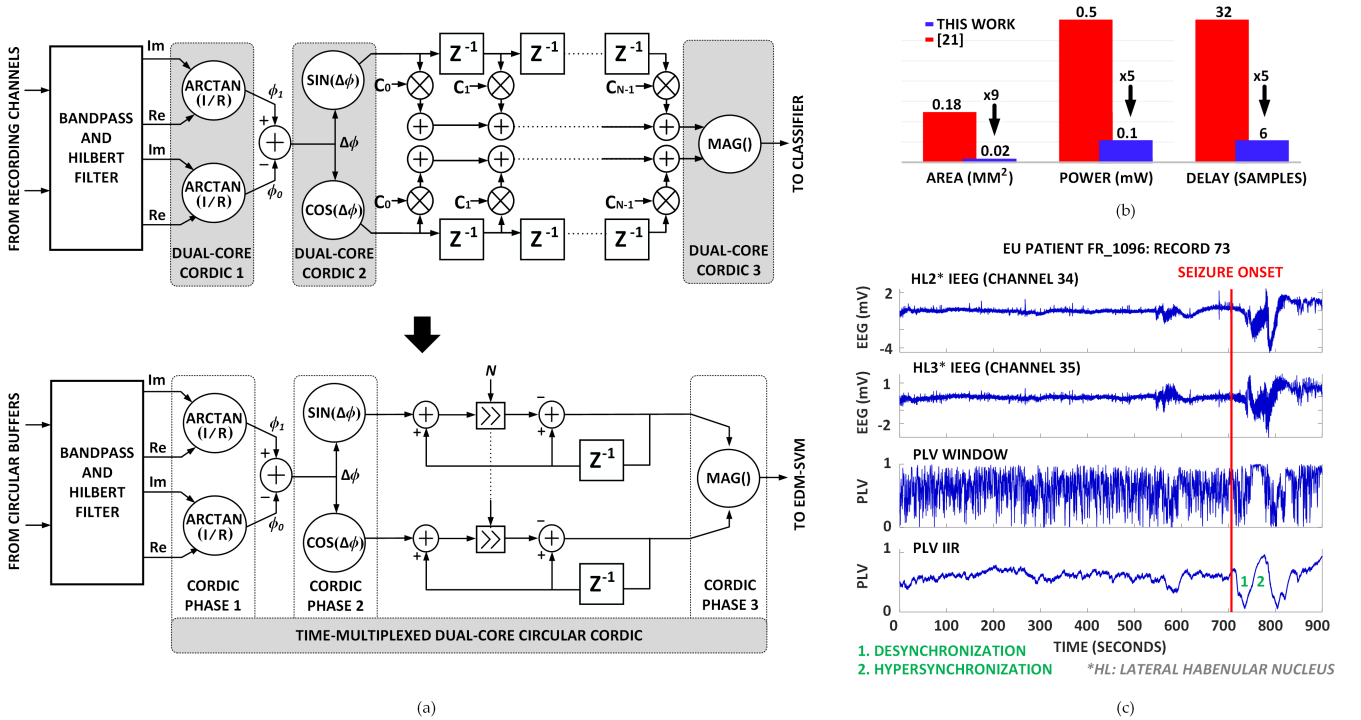

(a) (b) (c)

Fig. 4. (a) Phase locking value architecture compared to [25]. FIR based moving average filters are replaced with an IIR approximation. As only two CORDIC cores are used at a given time, time-multiplexing further reduces area. (b) Power, performance and area improvements with the implemented optimizations. (c) PLV biomarker extraction example from the EU epilepsy database showing a desynchronization at the onset of a seizure (1) followed by hypersynchronization (2) in the 7 Hz frequency band between 2 channels in the seizure onset zone extracted using a bit-accurate fixed-point MATLAB model.

to integrate functional brain regions and transfer information from global brain networks operating at behavioral timescales, to local high-frequency cortical processing [27]. In EEG, this mechanism manifests itself in local field potentials that resemble amplitude modulation in communication systems, where a neural signal can be described as:

$$x(t) = [1 + M.cos(2\pi f_{LF}(t) + \phi)] \cdot [A.sin(2\pi f_{HF}(t))] \quad (5)$$

Where f_{LF} is the low frequency modulating component, f_{HF} is the high frequency component whose amplitude is modulated, M is the modulation index and A is the base amplitude of the high frequency activity.

Elevated CFC between pathological high frequency oscillations and low frequency activity is found in the seizure-onset zone of epilepsy patients when compared to normal brain regions [28]. Abnormal PAC has also been found in the primary motor cortex of patients with Parkinson's Disease and a reduction has been shown in neural stimulation treatments which have alleviated symptoms [29]. These findings suggest that CFC could serve as a useful feedback measure in closed-loop neuromodulation devices.

The VLSI implementation of the approach presented in [30] allows for the use of two key metrics to enable a tradeoff between low-power, low-latency and high-precision. Two measures of CFC, the mean vector length modulation index (MVL-MI) and the cross-frequency phase locking value (CF-PLV) will be outlined in the following subsections.

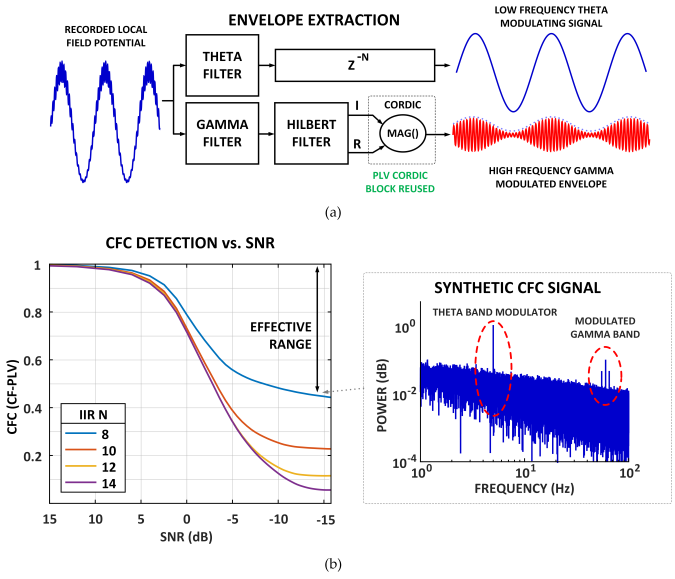

(a) (b)

Fig. 5. (a) The cross-frequency coupling envelope extractor isolates neuronal sub-population activity (red) which is modulated by global low-frequency oscillations (blue). (b) The noise tolerance of the feature extractor is characterized using a synthetic CFC signal in which theta-gamma coupling is corrupted using increasing 1/f noise.

1) Modulation Signal Extraction: Before calculating a CFC metric, it is first necessary to extract the low-frequency phase-modulating signal band, $f_{LF}(t)$, and high-frequency amplitude-modulated signal band, $f_{HF}(t)$ (Fig. 5). The raw

EEG signal, $x(t)$, is bandpass filtered using the previously introduced configurable-order FIR filter to extract both $f_{LF}(t)$ and $f_{HF}(t)$. After the modulated high-gamma component has been isolated, the amplitude envelope time series can be computed. An analytic signal is created using an FIR Hilbert filter and the dual-core CORDIC accelerator is then used to take the magnitude of the complex vector. This process extracts the envelope time series, $f_A(t)$, from $f_{HF}(t)$.

2) Mean Vector Length Modulation Index (MVL-MI) : The mean vector length modulation index (MVL-MI) finds a relationship between the instantaneous phases of $f_{LF}(t)$ and $f_A(t)$. This is achieved by building a complex-valued time series with a phase of $\phi_{LF}(t)$ and amplitude which is scaled by $f_A(t)$ [31]. This can be expressed as:

$$m(t) = |f_A(t)e^{i\phi_{LF}(t)}| \tag{6}$$

Where the phase-amplitude coupling measure, $m(t)$, is extracted from the time series defined in the complex plane. Each amplitude sample is represented by the length of the complex vector, whereas the phase of the modulating signal during that same sample is represented by the vector angle. When phase-amplitude coupling is not present, a uniform circular density of vector points is symmetric around zero. However, if the $f_{LF}(t)$ phase is modulating the high-frequency amplitude, the $f_{HF}(t)$ envelope is higher at certain phases. A measure of CFC can thus be quantified by taking the magnitude of the average complex vector [30].

3) Cross-Frequency Phase Locking Value (CF-PLV) : The cross-frequency phase locking value (CF-PLV) enables the detection of synchrony between the phase of the low frequency modulating signal $\phi_{fP(t)}$, and the phase of the envelope extracted from the high frequency modulated signal $\phi_{fA(t)}$. The PLV accelerator outlined in Section IV-B can be re-used, were the phase difference between both modulating and modulated signals is calculated as:

$$\Delta\phi(t) = \phi_{fA(t)} - \phi_{fP(t)} \tag{7}$$

As in between-channel PLV, the magnitude of an average vector is used as a measure of CFC between the phases of the modulating low-frequency signal and the envelope of the modulated high-frequency signal. If the average $\Delta\phi_i$ is 0, both $f_{LF}(t)$ and $f_{HF}(t)$ are phase-locked and CFC is present.

The ability of the hardware implementation presented in this work to detect CFC in the presence of noise is illustrated in Fig. 5. This is characterized with bit-accurate fixed-point stimulations using a synthetic CFC signal in which theta-gamma coupling is corrupted using increasing 1/f noise. It is noted that increasing values of N in the IIR filter approximation (Equation 3) increases the effective range of the output at the expense of decreased temporal resolution for transient changes in CFC.

D. Feature Extraction Summary

The feature extraction accelerators reduce the pre-classification workload which would otherwise need to be performed by the on-chip MSP430 CPU. A summary of this improvement in processing efficiency can be seen in Fig. 6. The output of these accelerators are normalized and sent to the proceeding machine learning subsystem for classification.

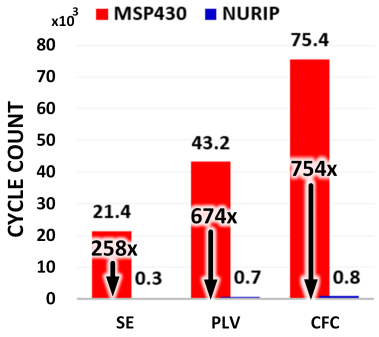

Fig. 6. Feature extraction on NURIP v.s. MSP430. Simulations using worst-case NURIP latency, MSP430 using the math.h library and with a dedicated hardware multiplier. The majority of the MSP430 overhead can be attributed to trigonometric functions which are accelerated in NURIP with dedicated CORDIC blocks.

TABLE I
NUMBER OF EXTRACTIONS (f_{clk} = 10MHz, f_{samp} = 256Hz)

	Throughput (# extractions)		Latency (cycles)	
FE	**Best Case**	**Worst Case**	**Best Case**	**Worst Case**
SE	151	151	258	258
PLV	241	57	162	674
CFC	161	51	242	754

The overall number of extractions which can be performed online is outlined in Table I. The largest contribution to the cycle count of these features is the required FIR filtering. Furthermore, the use of decimation filtering before feature extraction increases throughput by reducing the processing required. Such configurations can be determined on a per-patient basis to allow design tradeoffs based on recording channel noise and power-dissipation constraints.

V. Low-power Brain State Classification

For a given brain state, the biomarkers outlined in Section IV are expressed with high variability from patient to patient and change with the underlying physiology over time [32]. Data-driven approaches can be used to create models based on recorded iEEG data rather than using manual thresholds of feature values defined by a clinician. These models can then be used to personalize an implanted closed-loop medical device to accurately detect a patient's seizures.

While machine learning techniques based on deep-learning have attracted significant attention in recent times due to compelling performance in many classification tasks [33], their efficacy comes at a cost. In the context of seizure prediction, a recent study compares an LSTM/CNN deep learning approach with an SVM approach similar to the one presented in this work [34]. The authors demonstrate that deep learning and SVMs show comparable performance, but the computational complexity of the former is several orders of magnitudes higher (Table II). This efficiency renders the SVM employed in NURIP particularly suited to low-power

TABLE II
DEEP LEARNING V.S. SVM FOR LOW-POWER APPLICATIONS

Classifier	Sensitivity	FDR	MACs	Nonlinear Ops.
RCNN [34]†	95.74%	0.12	440776064	52289
SVM [35]*	96.09%	0.14	123906	384
Difference	**-0.35%**	**-0.02**	**x3557**	**x136**

† CNN: {3x3 conv. layer: 4, 2x2 pool. layer: 2}, LSTM: {128 hidden-unit: 60}, FCL: {60 hidden-unit: 1}

* Radial Basis Function Kernel, 321 support vectors, 384 features

seizure detection devices [6], [7]. A further key distinction between this work and deep-learning methods is that hand tuned features are employed with the SVM based on insights from domain knowledge. Deep learning approaches depart from such domain understanding, in favor of data-driven learning of features using complex models.

A. Exponentially Decaying Memory

Seizure onset biomarkers are subtle and can occur minutes before a clinical onset. Processing data on this timescale is a key challenge for implantable devices due to limited on-device memory. To capture the temporal evolution of biomarkers such as signal energy, conventional methods use a windowing approach where contiguous time epochs are concatenated to form a feature vector to be classified. The differences between these windows can be used to learn the characteristics of events such as the onset of a seizure [6], [32].

The use of windowing for time series classification has three main limitations: 1) The required device memory scales linearly with the history to be used for classification (Fig. 8). 2) As the test vectors are generated once per time epoch, the minimum detection latency is the time required to generate a window (typically multiple seconds) [6], [7]. 3) As temporal biomarkers are patient specific, so too should be the window length. The ability to learn appropriate timescales is therefore desirable to maximize the classifiers performance.

NURIP uses an exponentially decaying memory approach to represent complex temporal relationships for efficient classification and to addresses the outlined challenges associated with windowing. Rather than accumulating and concatenating fixed windows, a continuous sampling recursive window can be defined by:

$$EDM[t] = EDM[t-1] - \lambda(EDM[t-1] - x[t]) \quad (8)$$

Here, a new input $x[t]$ is incorporated based on a set learning rate, and the existing memory of a feature is degraded according to the decay rate, λ. Where:

$$\lambda = \frac{1}{2^{\alpha}}\ , 1 < \alpha < 16 \quad (9)$$

When the decay coefficient, λ, is constrained using the reciprocal of a power-of-two, the EDM update can efficiently be performed using shift and add operations. Each EDM has an approximate effective time window of $2^{\alpha}/Fs$ (Fig. 8).

The EDM's efficiency allows multiple units to be used in parallel to capture multiple timescales simultaneously (Fig. 7).

TABLE III
TEMPORAL STATE TABLE: QUANTIZED 2-EDM EXAMPLE

Temporal State Table		
State	**EDM 1**	**EDM 2**
1	LOW	LOW
2	AVG	LOW
3	LOW	AVG
4	HIGH	LOW
5	AVG	AVG
6	LOW	HIGH
7	HIGH	AVG
8	AVG	HIGH
9	HIGH	HIGH

This approach enables a classifier to process a feature's history across multiple timescales to learn temporal relationships. Taking quantized values as an example, Table III illustrates the use of two parallel EDMs to encode feature values in relative states, where an individual EDM state, $EDM \in \{LOW, AVG, HIGH\}$. In practice, the state space is much larger as each element has a resolution determined by the effective number of bits in the EDM.

B. EDM Support Vector Machine

The Support vector machine (SVM) is most commonly used as a supervised learning model for classification tasks of two or more classes. A data point is viewed as an N-dimensional vector and the fundamental idea is to find an N-dimensional hyperplane that can separate two groups of input data points which should also be mapped to the same high-dimensional space as the hyperplane [36].

The original algorithm requires a similar number of examples in each class to prevent classifier bias, but in the case of seizure detection, ictal activity is rare and prone to labeling errors. The SVM utilized in this work uses a semi-supervised one-class approach which has been proposed for datasets with such constraints [37]. It can be viewed as a regular two-class SVM where the training data is taken as one class, and the origin is taken as the only member of the second class. Training is performed using only interictal data which is mapped to the kernel space and separated from the origin by a hyperplane with maximum margin. The kernel, K, is implemented here using the Radial Basis Function (RBF):

$$K(\mathbf{sv}, \mathbf{x}) = e^{-\gamma\|\mathbf{sv}-\mathbf{x}\|^2} \quad (10)$$

Where sv are the support vectors used to construct the hyperplane, x is the extracted feature vector and γ is the inverse of the standard deviation of the RBF, or Gaussian function. Intuitively, the gamma parameter defines how far the influence of a single support vector reaches. The exponential function is computed using a shared hyperbolic CORDIC core. The implemented SVM accelerator allows the selection of linear, polynomial and RBF kernels to trade off between performance, energy and memory usage [6].

The combination of the EDM mechanism with SVM classification enables effective low-power time series classification. As the EDM is updated every sample, inference to

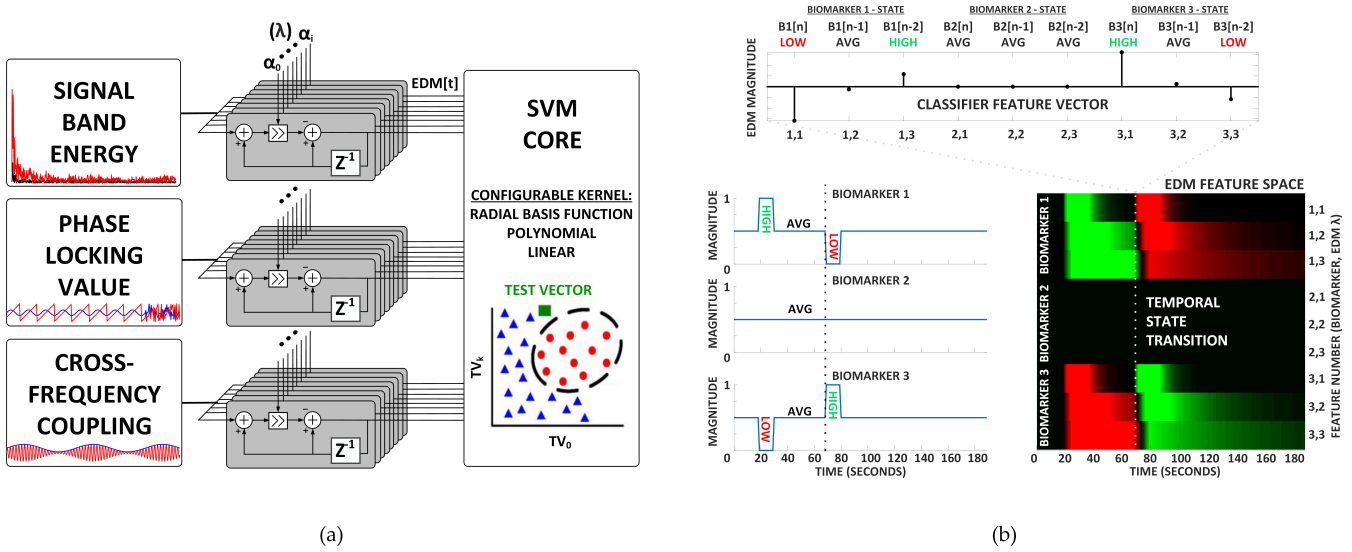

(a) (b)

Fig. 7. (a) The EDM's efficiency allows multiple biomarker timescales to be processed in parallel to provide a multi-timescale feature for the SVM core. (b) An array of EDMs with varying λ values encode relative temporal states in a feature for classification.

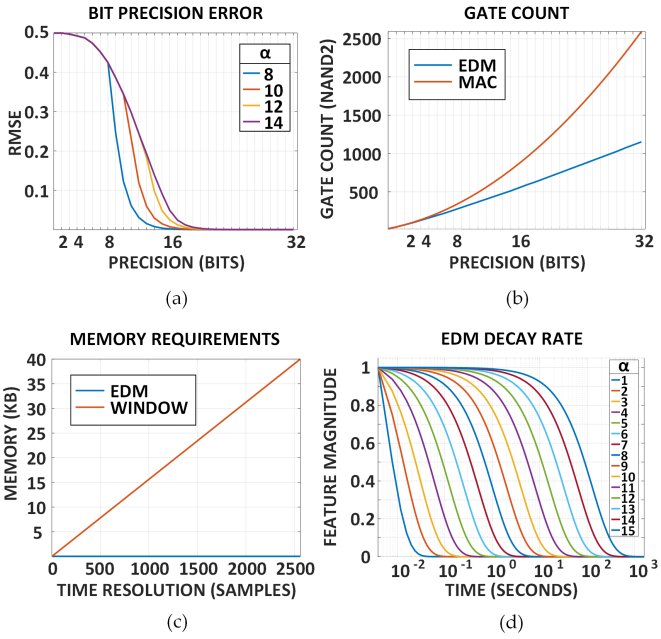

(a) (b) (c) (d)

Fig. 8. (a) The EDM feature bit precision requirement is characterized using Gaussian white noise by comparing the RMSE between an ideal MATLAB model and post-layout simulated EDMs synthesized with varying bit precision. (b) The gate count per bit of the EDM is lower and scales linearly compared to MAC-based window features. (c) Memory requirements scale linearly with time resolution using window features and are fixed when using EDM. (d) EDM decay response (assumes a sample rate of 256 Hz and 16-bit resolution).

be performed continuously rather than only when a window has been processed. Critically, the mechanism allows for the retention of biomarkers over time periods which are infeasible for windowing-based methods, where device memory requirements scale linearly with the number of samples. Furthermore, the efficiency of this mechanism allows for the combination of N long-term and short-term memory decay rate EDMs to enable the learning of complex temporal relationships as shown in Fig. 7.

VI. Classifier Performance

The processor has been verified and validated using neural recording data available in the EU Epilepsy database [15]. This database contains intracranial recordings from 30 patients with an average continuous recording time of 150 hours per patient. A review of the results from early seizure detection approaches using this dataset can be found in [38].

Due to the inherent class imbalance problem associated with seizure data (with few ictal examples compared to interictal data), sensitivity and false detection rate (FDR) measures are generally used in the seizure detection literature [35]. Sensitivity measures the proportion of real seizures that were correctly identified by a classifier while the false detection rate indicates the number of false alarms raised by a detector per hour of recording. NURIP's ability to detect clinically relevant brain states was evaluated using 500 hours of data from four patients in the EU Epilepsy database (Table IV).

Data was first downsampled to 256 Hz and 16 electrodes were chosen on a per patient basis based on their proximity to the seizure onset zone. The feature extraction uses five signal energy spectral bands, phase locking values between channels in the theta band and cross-frequency coupling between the theta and gamma frequency bands. EDM decay coefficients, of 8, 10, 12, and 14 are used (approximately 1, 4, 16, and 64 second effective time windows, respectively). A one-class SVM was trained using a radial basis function kernel. To estimate the classifier performance, a leave-one-record-out cross-validation scheme is employed as in [35].

Fig. 10 compares the performance of NURIP in terms of both sensitivity and FDR to the state of the art methods in [38] which are validated using the EU epilepsy database. A total of 43 of 44 clinical seizures analyzed were detected by the processor with a mean (SD) false positive rate of 0.185

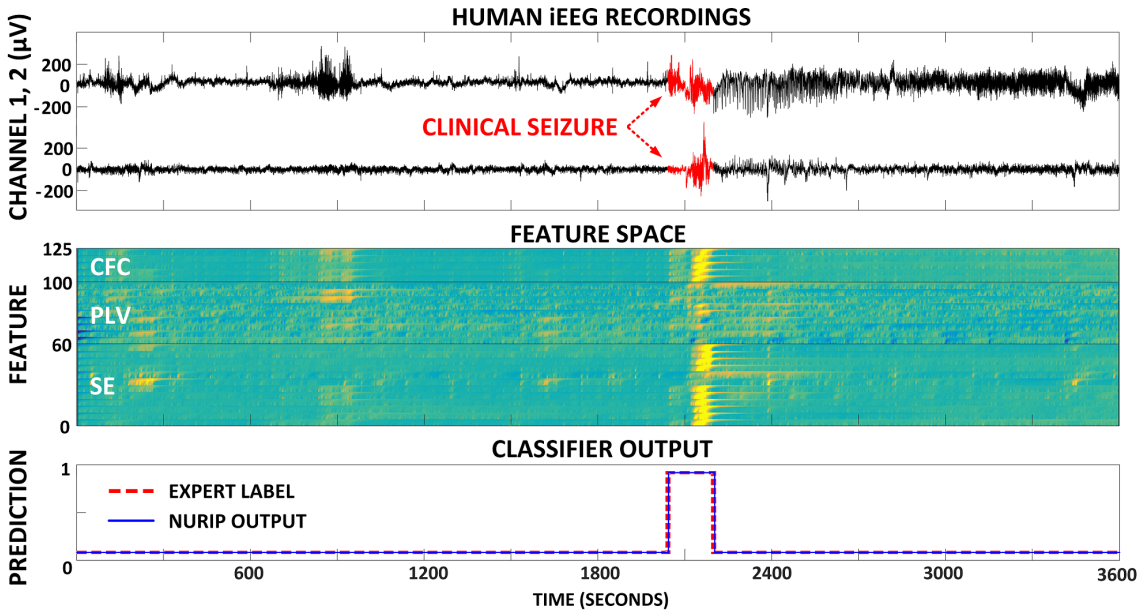

Fig. 9. Experimentally computed feature space and EDM-SVM classifier output in 1 hour of offline human iEEG recordings from the EU epilepsy database [15]. The extracted feature space used for classification consists of 125 dimensions derived through offline feature selection and is constrained to < 200 support vectors by the on-chip SRAM.

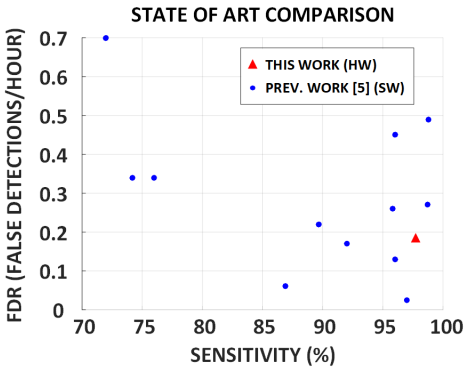

Fig. 10. Performance comparison of NURIP with the state of the art software (SW) methods in [38] in terms of both sensitivity and FDR.

TABLE IV
NURIP EU EPILEPSY DATABASE PERFORMANCE

EU Patient	Sex	Age	Hours Analyzed	Clinical Seizures	Seizures Detected	Sensitivity (%)	FDR
EU1096	F	32	147	7	7	100	0.16
EU442	M	21	118	8	8	100	0.11
EU548	M	17	129	17	16	94.10	0.22
EU1125	F	11	108	12	12	100	0.25
Sum	**2M/2F**	**-**	**502**	**44**	**43**	**-**	**-**
Mean (STD)	**-**	**20**	**125.5**	**-**	**-**	**97.7 (0.02)**	**0.185 (0.05)**

(0.05) per hour. The resulting sensitivity is greater than 84% of the existing work and the false detection rate is lower than 64% of those referenced. It should be noted that the classifiers included for comparison are based on software implementations with higher processing requirements.

VII. Programmable Charge-Balanced Neurostimulation Waveform Synthesizer

Upon the detection of a pathological brain state, an electrical stimulus can be applied to suppress symptoms. Three main challenges arise when implementing stimulation strategies for neuromodulation devices; which stimulation parameters are necessary to induce the desired neural activity, and hence produce the desired effects? How can we minimize the power required to achieve a given effect? And how can we ensure our stimulation parameters are safe for chronic use?

It has been demonstrated the use of stimulus waveforms with a net direct current component increases the probability of tissue and electrode damage [39]. Most neural stimulators today deliver charge-balanced bi-phasic rectangular current pulses, where the first (cathodic) phase excites the nerve fiber and the second (anodic) phase provides charge balancing. The rectangular waveform is widely used for its simplicity and ease of generation with a simple current source. However, it has been shown that arbitrary waveforms can induce complex neural activity. Indeed, this is the mechanism behind modern cochlear implants, where non-bi-phasic electrical stimulation has been used to control auditory neurons to convey meaningful information to the brain [40]. Approaches have been developed to learn the optimal stimulation parameters to trade off between selectivity, reduced power consumption and waveform safety [11]. The ability to programmatically control these parameters on a device could allow online waveform adaptation based on such closed-loop control techniques.

NURIP integrates a digitally charge-balanced neurostimula-

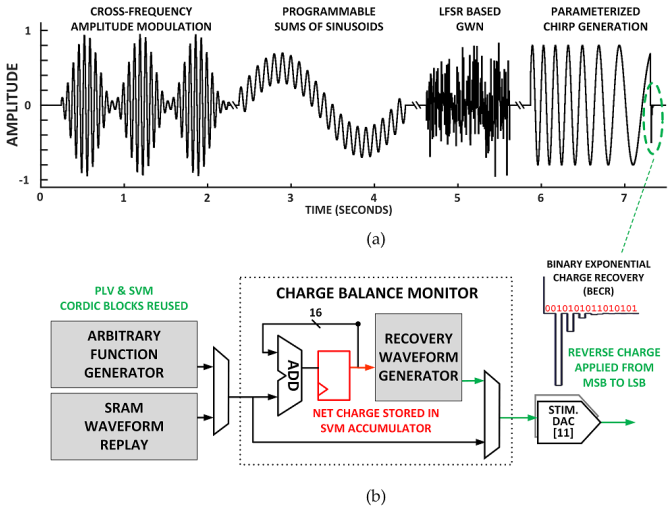

Fig. 11. (a) Experimentally measured arbitrary waveform generation and charge balancing. (b) The SVM MAC accumulator is reused to maintain the net charge sent to an analog front end neural stimulator. When the stimulus is terminated or when safety limits have been exceeded, the inverse of the sum is iteratively applied moving from MSB to LSB.

tion waveform synthesizer which is demonstrated in Fig. 11 for arbitrary waveform and function generation (AWFG). Function generation up to 15kHz is efficiently implemented with the re-use of the dual-core CORDIC and MAC blocks to support the generation of sums or products of sinusoids. 3MHz arbitrary waveform replay is supported by streaming samples from on-chip SRAM.

To mitigate the issue of charge imbalanced stimulation, Ref. [41] proposes the pulsating voltage transcranial electrical stimulator (PVTES) to adapt the number of stimulation pulses with respect to skin-electrode impedance variation. However, this approach is limited to bi-phasic pulse stimulation. A digital charge-balancing technique is implemented in this work to support the use of arbitrary waveforms. The system MAC logic is re-used to store the net charge sent to an analog front end neural stimulator and hence monitor the stimulus to ensure safe limits are not exceeded. An exponential charge recovery phase has been demonstrated to safely reduce such imbalances when compared to sudden terminations [39]. This approach is efficiently implemented in NURIP using binary exponential charge recovery (BECR) where the inverse of the charge monitoring register is iteratively applied from MSB to LSB. While BECR ensures that the digital values sent to analog DACs do not exceed safe limits, the charge monitoring register should be adjusted with feedback from the analog domain to compensate for stimulator nonidealities and varying electrode impedances.

VIII. VLSI Implementation and Measurement

The processor is implemented in a 0.13-μm RF CMOS process from IBM as shown in the micrograph in Fig. 12. The design utilizes an area of 2.55x1.3mm with a logic size of 509k NAND2 equivalent gates using an ARM RVT standard cell library. A power consumption of 674.4 μW

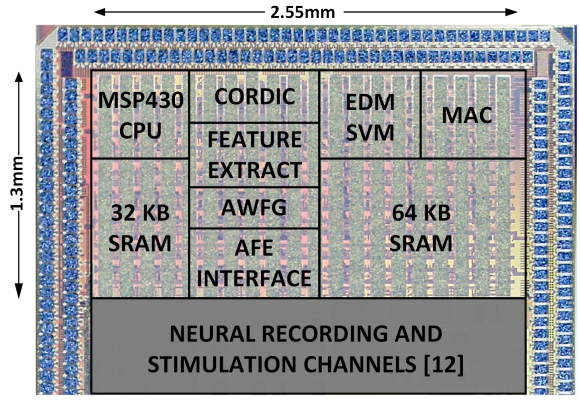

Fig. 12. NURIP chip micrograph with major blocks labeled.

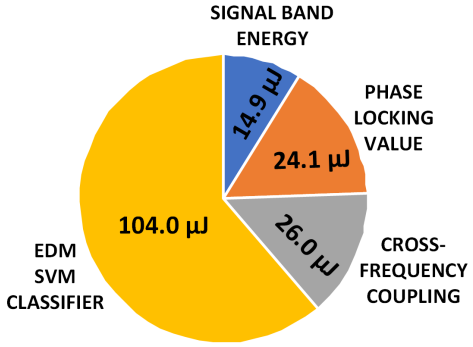

Fig. 13. Energy usage breakdown for a typical classification based on experimentally measured power consumption.

with a nominal voltage of 1.2V was measured using a EDM-SVM classification rate of 4Hz with a radial basis function kernel, and continuous feature extraction with an operational frequency of 10 MHz. As outlined in the power breakdown in Fig. 13, convolutions for feature extraction and classification represent the majority of the devices activity.

Functional verification was performed using an external debug interface to on-chip logic for data streaming and control. Samples from the EU database are streamed to the preprocessor via a test FPGA for analysis and generated digital waveform values are accessed via memory mapped registers for visualization as shown in Fig. 11.

The SoC is compared with the state of the art in Table V. NURIP implements the broadest range of feature extractors at a cost of power consumption, but its classification performance is among the highest demonstrated using the EU Epilepsy Database as shown in Fig. 10 (which includes non-implantable approaches).

IX. Conclusion

In this work, NURIP integrates accurate brain state classification for patient-specific seizure detection with neuromodulation waveform generation for precise simulation and contingent seizure abortion with a processing power consumption of 674.4 μW. The on-chip autoencoder structure for signal conditioning and dimensionality reduction is an approach that

TABLE V
COMPARISON TABLE

	JSSC'13 [5]	JSSC'11 [13]	JSSC'14 [14]	JSSC'15 [7]	JSSC'17 [4]	THIS WORK
TECHNOLOGY (μm)	0.13	0.13	0.18	0.18	0.13	0.13
SUPPLY VOLTAGE	0.55-1.2	0.5-1.0	1.8	1.0	1.2	1.2
ANALOG FRONT END	NO	NO	YES	YES	YES	YES
SRAM (kB)	64	128	-	64	-	96
FEATURE EXTRACTION	CPU	SE, FFT	FFT, Entropy	SE	PLV	PLV, CFC, SE, CPU
CLASSIFIER	SVM	CPU	LLS	D2A-LSVM	Threshold	EDM-SVM
ENERGY/CLASS. (μJ)	273	-	77.91	2.73	-	168.6
SAMPLE MEMORY (s)	6	0.125	0.375	3	0.24	multi-minute (EDM)
WIN. LATENCY (s)	2	-	0.8	1	-	0
SENSITIVITY (%)	100*	-	92†	95.7*	75†	97.7◇
FDR	0.05*	-	-	0.27*	1.0†	0.185◇
WAVEFORM SYNTHESIS	-	-	Bi-phasic	Bi-phasic	-	AWFG
CHARGE BALANCING	-	-	-	PVTES	-	BECR

Evaluated using: *MIT-CHB Database, †Local data, ◇EU Epilepsy Database

could greatly reduce device computational requirements as we scale towards higher channel counts.

The SE, PLV and CFC array of feature extractors combined with on-chip machine learning allows the classification other brain states to be explored, such as those found in Parkinson's disease. The EDM approach to time series classification has been demonstrated to efficiently encode temporal relationships in biomarkers for classification and overcome the high memory requirement associated with windowing. This approach could be expanded to other applications where long-term dependencies should be considered.

On-chip neuromodulation waveform synthesis using the approaches presented here enables future adaptive stimulation paradigms to be explored for more precise control of neural activity. The BECR charge balancing solution in combination with analog monitoring will ensure that such paradigms do not increase the risk of electrode and tissue damage in patients.

Integrating neural signal acquisition, signal processing, machine learning and neuromodulation waveform generation on single device solutions has the potential to revolutionize the quality treatment for those with neurological disorders.

ACKNOWLEDGMENT

The authors thank the Canadian Microelectronics Corporation (CMC) for fabrication, and the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding. The authors also thank M. R. Pazhouhandeh for his assistance with top-level verification and PCB design, and Jintao Zhang for valuable discussions.

REFERENCES

[1] G. Buzski, C. A. Anastassiou, and C. Koch, "The origin of extracellular fields and currents EEG, ECoG, LFP and spikes," *Nat Rev Neurosci*, vol. 13, no. 6, pp. 407–420, May 2012. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4907333/

[2] C. B. Saper, P. M. Fuller, N. P. Pedersen, J. Lu, and T. E. Scammell, "Sleep State Switching," *Neuron*, vol. 68, no. 6, pp. 1023–1042, Dec. 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0896627310009748

[3] G. K. Bergey, M. J. Morrell, E. M. Mizrahi, A. Goldman, D. King-Stephens, D. Nair, S. Srinivasan, B. Jobst, R. E. Gross, D. C. Shields, G. Barkley, V. Salanova, P. Olejniczak, A. Cole, S. S. Cash, K. Noe, R. Wharen, G. Worrell, A. M. Murro, J. Edwards, M. Duchowny, D. Spencer, M. Smith, E. Geller, R. Gwinn, C. Skidmore, S. Eisenschenk, M. Berg, C. Heck, P. Van Ness, N. Fountain, P. Rutecki, A. Massey, C. O'Donovan, D. Labar, R. B. Duckrow, L. J. Hirsch, T. Courtney, F. T. Sun, and C. G. Seale, "Long-term treatment with responsive brain stimulation in adults with refractory partial seizures," *Neurology*, vol. 84, no. 8, pp. 810–817, Feb. 2015. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4339127/

[4] H. Kassiri, M. T. Salam, M. R. Pazhouhandeh, N. Soltani, J. L. Perez Velazquez, P. Carlen, and R. Genov, "Rail-to-Rail-Input Dual-Radio 64-Channel Closed-Loop Neurostimulator," *IEEE Journal of Solid-State Circuits*, pp. 1–18, 2017. [Online]. Available: http://ieeexplore.ieee.org/document/8068946/

[5] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. P. Chandrakasan, "A Micro-Power EEG Acquisition SoC With Integrated Feature Extraction Processor for a Chronic Seizure Detection System," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 4, pp. 804–816, Apr. 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5437484/

[6] K. H. Lee and N. Verma, "A Low-Power Processor With Configurable Embedded Machine-Learning Accelerators for High-Order and Adaptive Analysis of Medical-Sensor Signals," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 7, pp. 1625–1637, Jul. 2013. [Online]. Available: http://ieeexplore.ieee.org/document/6493458/

[7] M. A. Bin Altaf, C. Zhang, and J. Yoo, "A 16-Channel Patient-Specific Seizure Onset and Termination Detection SoC With Impedance-Adaptive Transcranial Electrical Stimulator," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 11, pp. 2728–2740, Nov. 2015. [Online]. Available: http://ieeexplore.ieee.org/document/7302516/

[8] M. T. Salam, H. Kassiri, R. Genov, and J. L. Perez Velazquez, "Rapid brief feedback intracerebral stimulation based on real-time desynchronization detection preceding seizures stops the generation of convulsive paroxysms," *Epilepsia*, vol. 56, no. 8, pp. 1227–1238, Aug. 2015. [Online]. Available: http://doi.wiley.com/10.1111/epi.13064

[9] M. Sahin and Y. Tie, "Non-rectangular waveforms for neural stimulation with practical electrodes," *J Neural Eng*, vol. 4, no. 3, pp. 227–233, Sep. 2007.

[10] A. Wongsarnpigoon and W. M. Grill, "Energy-efficient waveform shapes for neural stimulation revealed with genetic algorithm," *J Neural Eng*, vol. 7, no. 4, p. 046009, Aug. 2010. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2925408/

[11] J. Pineau and a. et, "Treating epilepsy via adaptive neurostimulation: a reinforcement learning approach," *Int J Neural Syst*, vol. 19, no. 4, pp. 227–240, Aug. 2009. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4884089/

[12] G. O'Leary, M. R. Pazhouhandeh, M. Chang, D. Groppe, T. A. Valiante, N. Verma, and R. Genov, "A recursive-memory brain-state classifier with 32-channel track-and-zoom D2s ADCs and Charge-Balanced Programmable Waveform Neurostimulators," in *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*, Feb. 2018, pp. 296–298.

[13] J. Kwong and A. P. Chandrakasan, "An Energy-Efficient Biomedical Signal Processing Platform," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 7, pp. 1742–1753, Jul. 2011.

[14] W.-M. Chen, H. Chiueh, T.-J. Chen, C.-L. Ho, C. Jeng, M.-D. Ker, C.-Y. Lin, Y.-C. Huang, C.-W. Chou, T.-Y. Fan, M.-S. Cheng, Y.-L. Hsin, S.-F. Liang, Y.-L. Wang, F.-Z. Shaw, Y.-H. Huang, C.-H. Yang, and C.-Y. Wu, "A Fully Integrated 8-Channel Closed-Loop Neural-Prosthetic CMOS SoC for Real-Time Epileptic Seizure Control," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 232–247, Jan. 2014. [Online]. Available: http://ieeexplore.ieee.org/document/6637111/

[15] M. Ihle, H. Feldwisch-Drentrup, C. A. Teixeira, A. Witon, B. Schelter, J. Timmer, and A. Schulze-Bonhage, "EPILEPSIAE A European epilepsy database," *Computer Methods and Programs in Biomedicine*, vol. 106, no. 3, pp. 127–138, Jun. 2012. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0169260710002221

[16] K. A. Ludwig, R. M. Miriani, N. B. Langhals, M. D. Joseph, D. J. Anderson, and D. R. Kipke, "Using a Common Average Reference to Improve Cortical Neuron Recordings From Microelectrode Arrays," *J Neurophysiol*, vol. 101, no. 3, pp. 1679–1689, Mar. 2009. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2666412/

[17] C. M. Lopez, S. Mitra, J. Putzeys, B. Raducanu, M. Ballini, A. Andrei, S. Severi, M. Welkenhuysen, C. Van Hoof, S. Musa, and R. F. Yazicioglu, "A 966-Electrode Neural Probe with 384 Configurable Channels in 0.13m SOI CMOS," 2016. [Online]. Available: http://eprints.gla.ac.uk/118034/

[18] E. Plaut, "From Principal Subspaces to Principal Components with Linear Autoencoders," *arXiv:1804.10253 [cs, stat]*, Apr. 2018, arXiv: 1804.10253. [Online]. Available: http://arxiv.org/abs/1804.10253

[19] N. Soltani, M. S. Aliroteh, M. T. Salam, J. L. P. Velazquez, and R. Genov, "Low-Radiation Cellular Inductive Powering of Rodent Wireless Brain Interfaces: Methodology and Design Guide," 2016. [Online]. Available: http://ieeexplore.ieee.org/xpls/absall.jsp?arnumber=7426419

[20] P. D. Wolf, "Thermal Considerations for the Design of an Implanted Cortical BrainMachine Interface (BMI)," in *Indwelling Neural Implants: Strategies for Contending with the In Vivo Environment*, W. M. Reichert, Ed. Boca Raton (FL): CRC Press/Taylor & Francis, 2008. [Online]. Available: http://www.ncbi.nlm.nih.gov/books/NBK3932/

[21] K. Abdelhalim, H. M. Jafari, L. Kokarovtseva, J. L. P. Velazquez, and R. Genov, "64-Channel UWB Wireless Neural Vector Analyzer SOC With a Closed-Loop Phase Synchrony-Triggered Neurostimulator," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 10, pp. 2494–2510, Oct. 2013.

[22] "Oppenheim, Schafer & Buck, Discrete-Time Signal Processing | Pearson." [Online]. Available: https://www.pearson.com/us/higher-education/product/Oppenheim-Discrete-Time-Signal-Processing-2nd-Edition/9780137549207.html

[23] J.-P. Lachaux, E. Rodriguez, J. Martinerie, F. J. Varela, and others, "Measuring phase synchrony in brain signals," *Human brain mapping*, vol. 8, no. 4, pp. 194–208, 1999. [Online]. Available: http://www.ma.utexas.edu/users/davis/reu/ch3/cwt/lachaux.pdf

[24] J. Volder, "The CORDIC computing technique," in *afips*. IEEE, 1899, p. 257.

[25] K. Abdelhalim, V. Smolyakov, and R. Genov, "Phase-synchronization early epileptic seizure detector VLSI architecture," *IEEE transactions on biomedical circuits and systems*, vol. 5, no. 5, pp. 430–438, 2011. [Online]. Available: http://ieeexplore.ieee.org/xpls/absall.jsp?arnumber=6046232

[26] R. T. Canolty and R. T. Knight, "The functional role of cross-frequency coupling," *Trends in Cognitive Sciences*, vol. 14, no. 11, pp. 506–515, Nov. 2010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1364661310002068

[27] A. B. L. Tort, R. Komorowski, H. Eichenbaum, and N. Kopell, "Measuring Phase-Amplitude Coupling Between Neuronal Oscillations of Different Frequencies," *Journal of Neurophysiology*, vol. 104, no. 2, pp. 1195–1210, Aug. 2010. [Online]. Available: http://jn.physiology.org/cgi/doi/10.1152/jn.00106.2010

[28] M. Guirgis, Y. Chinvarun, P. L. Carlen, and B. L. Bardakjian, "The role of delta-modulated high frequency oscillations in seizure state classification," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2013, pp. 6595–6598, 2013.

[29] C. de Hemptinne, N. C. Swann, J. L. Ostrem, E. S. Ryapolova-Webb, M. San Luciano, N. B. Galifianakis, and P. A. Starr, "Therapeutic deep brain stimulation reduces cortical phase-amplitude coupling in Parkinson's disease," *Nature Neuroscience*, vol. 18, no. 5, pp. 779–786, Apr. 2015. [Online]. Available: http://www.nature.com/doifinder/10.1038/nn.3997

[30] G. O'Leary, T. A. Valiante, and R. Genov, "Low-latency VLSI architecture for neural cross-frequency coupling analysis," in *Engineering in Medicine and Biology Society (EMBC), 2017 39th Annual International Conference of the IEEE*. IEEE, 2017, pp. 2247–2250.

[31] R. T. Canolty, E. Edwards, S. S. Dalal, M. Soltani, S. S. Nagarajan, H. E. Kirsch, M. S. Berger, N. M. Barbaro, and R. T. Knight, "High gamma power is phase-locked to theta oscillations in human neocortex," *science*, vol. 313, no. 5793, pp. 1626–1628, 2006. [Online]. Available: http://science.sciencemag.org/content/313/5793/1626.short

[32] G. O'Leary, A. O. Abraham, A. K. Kamath, D. Groppe, T. A. Valiante, and R. Genov, "Machine learning microserver for neuromodulation device training," in *2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, Oct. 2017, pp. 1–4.

[33] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015. [Online]. Available: http://www.nature.com/articles/nature14539

[34] P. Thodoroff, J. Pineau, and A. Lim, "Learning Robust Features using Deep Learning for Automatic Seizure Detection," *Machine Learning and Healthcare Conference (MLHC 2016), Los Angeles, CA*, Jul. 2016, arXiv: 1608.00220. [Online]. Available: http://arxiv.org/abs/1608.00220

[35] A. Shoeb and J. Guttag, "Application of Machine Learning to Epileptic Seizure Detection," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ser. ICML'10. USA: Omnipress, 2010, pp. 975–982. [Online]. Available: http://dl.acm.org/citation.cfm?id=3104322.3104446

[36] V. Vapnik, *The Nature of Statistical Learning Theory*, 2nd ed., ser. Information Science and Statistics. New York: Springer-Verlag, 2000. [Online]. Available: //www.springer.com/gp/book/9780387987804

[37] B. Schlkopf, R. C. Williamson, A. J. Smola, J. Shawe-Taylor, and J. C. Platt, "Support vector method for novelty detection," in *Advances in neural information processing systems*, 2000, pp. 582–588.

[38] M. Bandarabadi, J. Rasekhi, C. A. Teixeira, T. I. Netoff, K. K. Parhi, and A. Dourado, "Early Seizure Detection Using Neuronal Potential Similarity: A Generalized Low-Complexity and Robust Measure," *International Journal of Neural Systems*, vol. 25, no. 05, p. 1550019, Aug. 2015. [Online]. Available: http://www.worldscientific.com/doi/abs/10.1142/S0129065715500197

[39] D. R. Merrill, M. Bikson, and J. G. Jefferys, "Electrical stimulation of excitable tissue: design of efficacious and safe protocols," *Journal of Neuroscience Methods*, vol. 141, no. 2, pp. 171–198, Feb. 2005. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0165027004003826

[40] M. Yip, R. Jin, H. H. Nakajima, K. M. Stankovic, and A. P. Chandrakasan, "A Fully-Implantable Cochlear Implant SoC With Piezoelectric Middle-Ear Sensor and Arbitrary Waveform Neural Stimulation," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 1, pp. 214–229, Jan. 2015.

[41] M. A. B. Altaf and J. Yoo, "A 1.83 muJ/Classification, 8-Channel, Patient-Specific Epileptic Seizure Classification SoC Using a Non-Linear Support Vector Machine," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 10, no. 1, pp. 49–60, Feb. 2016.

Gerard O'Leary received the B.Eng. degree at the National University of Ireland, Galway, and the M.A.Sc. degree in electrical and computer engineering at the University of Toronto. He has held engineering positions at Analog Devices, The European Space Agency, and ARM. He is currently pursuing the Ph.D. degree in electrical and computer engineering and the collaborative specialization in biomedical engineering at the University of Toronto. His research is at the interface between neural and electronic systems, creating low-power integrated circuits which utilize signal processing and machine learning techniques to infer brain states and control activity using responsive neuromodulation.

David M. Groppe is a clinical data scientist who has spent the past 15 years developing novel statistical methods and software for mining human electrophysiology data. He has given multiple invited talks on analysis of electrophysiology data and one of his papers was named one of 10 must read papers in the influential textbook, An Introduction to the Event-Related Potential Technique. He is the lead developer of iELVis, a software package for analyzing data acquired from epilepsy surgery candidates, which is now in use at some of the worlds top epilepsy research centres (e.g., the Krembil Neuroscience Centre, Mayo Clinic, and Stanford). Using this package, he developed the first quantitative atlas of intracranially recorded human brain rhythms.

Taufik A. Valiante graduated from the MD/PhD program at the University of Toronto in 1997. In that year he began his residency in the Neurosurgery Training Program at the University of Toronto, and he became a Fellow of the Royal College of Physicians and Surgeons of Canada in 2003. Dr. Valiante then pursued an epilepsy neurosurgery fellowship at the University of Washington, Seattle, under the mentorship of Dr George Ojemann. In 2003, he was recruited back to the University of Toronto. He is the Director of the Surgical Epilepsy Program at the Krembil Neuroscience Center, Scientist at the Krembil Research Institute, and the Co-Director of CRANIA (the CenteR for Advancing Neurotechnological Innovation to Application). His neurosurgical staff appointment is at the Toronto Western Hospital, and he is currently Associate Professor of Neurosurgery, with cross-appointments in the Institute of Biomaterials and Biomedical Engineering, and Electrical and Computer Engineering, at the University of Toronto.

Naveen Verma received the B.A.Sc. degree in Electrical and Computer Engineering from the UBC, Vancouver, Canada in 2003, and the M.S. and Ph.D. degrees in Electrical Engineering from MIT in 2005 and 2009 respectively. Since July 2009 he has been with the Department of Electrical Engineering at Princeton University, where he is currently an Associate Professor. His research focuses on advanced sensing systems, exploring how systems for learning, inference, and action planning can be enhanced by algorithms that exploit new sensing and computing technologies. This includes research on large-area, flexible sensors, energy-efficient statistical-computing architectures and circuits, and machine-learning and statistical-signal-processing algorithms. Prof. Verma has served as a Distinguished Lecturer of the IEEE Solid-State Circuits Society, and currently serves on the technical program committees for ISSCC, VLSI Symp., DATE, and IEEE Signal-Processing Society (DISPS). Prof. Verma is recipient or co-recipient of the 2006 DAC/ISSCC Student Design Contest Award, 2008 ISSCC Jack Kilby Paper Award, 2012 Alfred Rheinstein Junior Faculty Award, 2013 NSF CAREER Award, 2013 Intel Early Career Award, 2013 Walter C. Johnson Prize for Teaching Excellence, 2013 VLSI Symp. Best Student Paper Award, 2014 AFOSR Young Investigator Award, 2015 Princeton Engineering Council Excellence in Teaching Award, and 2015 IEEE Trans. CPMT Best Paper Award.

Roman Genov received the B.S. degree in Electrical Engineering from Rochester Institute of Technology, NY in 1996 and the M.S.E. and Ph.D. degrees in Electrical and Computer Engineering from Johns Hopkins University, Baltimore, MD in 1998 and 2003 respectively. He is currently a Professor in the Department of Electrical and Computer Engineering at the University of Toronto, Canada, where he is a member of Electronics Group and Biomedical Engineering Group and the Director of Intelligent Sensory Microsystems Laboratory. Dr. Genovs research interests are primarily in analog integrated circuits and systems for energy-constrained biological, medical, and consumer sensory applications. Dr. Genov is a co-recipient of Jack Kilby Award for Outstanding Student Paper at IEEE International Solid-State Circuits Conference, NSERC Discovery Accelerator Award, Best Paper Award by IEEE Transactions on Biomedical Circuits and Systems, Best Paper Award at IEEE Biomedical Circuits and Systems Conference, Best Paper Award in Biomedical Systems and Best Student Paper Award at IEEE International Symposium on Circuits and Systems, Brian L. Barge Award for Excellence in Microsystems Integration, MEMSCAP Microsystems Design Award, DALSA Corporation Award for Excellence in Microsystems Innovation, and Canadian Institutes of Health Research Next Generation Award. He was a Technical Program Co-chair at IEEE Biomedical Circuits and Systems Conference. He was an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-II: EXPRESS BRIEFS and IEEE SIGNAL PROCESSING LETTERS, and a member of IEEE International Solid-State Circuits Conference International Program Committee. Currently he is an Associate Editor of IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS, and its Guest Editor of Special Issues on IEEE International Solid-State Circuits Conference. He is also a member of the European Solid-State Circuits Conference International Program Committee.