LETTER Communicated by Paolo Del Giudice

Dynamic State and Parameter Estimation Applied to Neuromorphic Systems

Emre Ozgur Neftci
emre@ini.phys.ethz.ch
Institute of Informatics, University of Zurich and ETH Zurich, Zurich CH-8057, Switzerland

Bryan Toth
btoth@physics.ucsd.edu
Department of Physics and Center for Theoretical Biological Physics, Scripps Institute of Oceanography, University of California, San Diego, La Jolla, CA 92093-0402, U.S.A.

Giacomo Indiveri
giacomo@ini.phys.ethz.ch
Institute of Informatics, University of Zurich and ETH Zurich, Zurich CH-8057, Switzerland

Henry D. I. Abarbanel
habarbanel@ucsd.edu
Department of Physics and Marine Physical Laboratory, Scripps Institute of Oceanography, University of California, San Diego, La Jolla, CA 92093-0402, U.S.A.

Neuroscientists often propose detailed computational models to probe the properties of the neural systems they study. With the advent of neuromorphic engineering, there is an increasing number of hardware electronic analogs of biological neural systems being proposed as well. However, for both biological and hardware systems, it is often difficult to estimate the parameters of the model so that they are meaningful to the experimental system under study, especially when these models involve a large number of states and parameters that cannot be simultaneously measured. We have developed a procedure to solve this problem in the context of interacting neural populations using a recently developed dynamic state and parameter estimation (DSPE) technique. This technique uses synchronization as a tool for dynamically coupling experimentally measured data to its corresponding model to determine its parameters and internal state variables. Typically experimental data are obtained from the biological neural system and the model is simulated in software; here we show that this technique is also efficient in validating proposed

network models for neuromorphic spike-based very large-scale integration (VLSI) chips and that it is able to systematically extract network parameters such as synaptic weights, time constants, and other variables that are not accessible by direct observation. Our results suggest that this method can become a very useful tool for model-based identification and configuration of neuromorphic multichip VLSI systems.

1 Introduction

Perhaps the most famous example of biophysical modeling of a neural system is the original experimental work of Hodgkin and Huxley (1952), which revealed, among many other aspects of neuron dynamics, the mechanisms for spike generation and propagation. To date, as Hodgkin and Huxley (1952) did, electrophysiologists still employ techniques such as current, voltage, and conductance clamping to determine the parameters of their model (Jolivet, Lewis, & Gerstner, 2004; Brette & Gerstner, 2005). Unfortunately, these techniques are difficult to apply systematically, especially when some of the variables cannot be directly observed. For this reason, researchers have developed a wide variety of techniques for determining model parameters from experimental data, which includes, for example maximum likelihood methods (Brillinger, 1988; Huys, Ahrens, & Paninski, 2006; Paninski, Pillow, & Simoncelli, 2004; Okatan, Wilson, & Brown, 2005) or metaheuristic methods (Rossant, Goodman, Platkiewicz, & Brette, 2010). However, there is still no common framework for fitting methods. This lack has recently led to the International Neuroinformatics Coordinating Facility (INCF) Quantitative Single-Neuron Modeling Competition (Naud, Berger, Gerstner, Bathellier, & Carandini, 2009). This competition, intended to bridge the gap between electrophysiologists and modelers, underlines the fact that a systematic method for model and parameter estimation is a necessity in neuroscience.

In this work, we propose a novel method for determining the properties of multiple interacting neural populations that is systematic and does not require the measurement of membrane potentials. The method proposed applies to both model validation and parameter estimation and is based on the recently proposed dynamic state and parameter estimation (DSPE) technique of Abarbanel, Creveling, Farsian, and Kostuk (2009) (see also Toth, Kostuk, Meliza, Margoliash, & Abarbanel, 2011). The DSPE technique builds on existing work that uses synchronization as a tool for parameter estimation and incorporates information from experimental observations in an efficient manner. It has been previously applied to a simulated standard Hodgkin and Huxley (H&H) neuron using recordings of the membrane potential and was able to accurately fit the 22 parameters of the model and estimate the three unobserved states (Abarbanel et al., 2009).

The parameter estimation in networks of H&H neurons is a much more difficult problem, because it introduces many nonlinear interactions at very different timescales (see Wang & Buzsáki, 1996, for such a model). We focus on a more tractable aspect of this problem, which is the determination of population-level parameters (Wilson & Cowan, 1972, 1973; Abbott, 1994). This level of detail in neural modeling is often described by neural masses (Freeman, 1975) and models the interactions between neural populations without studying the detailed spatial structure of the connections. The study of dynamics in neural masses is argued to be the relevant granularity for modeling purposes in neural systems (Freeman, 1975; Faugeras, Grimbert, & Slotine, 2008). This type of model is often employed to understand the neural mechanisms underlying electrophysiological measurements such as EEG or functional neural imaging data, obtained from fMRI or PET imaging methods (David & Friston, 2003; Freeman, Ahlfors, & Menon, 2009). In this letter, we focus on the parameter estimation for neuromorphic very large-scale integration (VLSI) neural networks, where this level of modeling is used for model-based identification and chip configuration purposes.

1.1 Application to Neuromorphic Systems. An increasing number of neuroscientists, computer scientists and engineers are now converging their efforts in designing VLSI neuromorphic devices and systems (Mead, 1989; Indiveri & Horiuchi, 2011). In particular, a large subset of the neuromorphic engineering field focuses on emulating the properties of biological neurons in VLSI circuits (Mead, 1989; Mahowald & Douglas, 1991; Indiveri et al., 2011). The goal of these emulations is to understand and exploit the physical principles of brainlike computation for novel computational technologies rather than to simulate its detailed biophysics on general-purpose digital computers.

Although there are several advantages to this approach (Douglas, Mahowald, & Mead, 1995; Indiveri & Horiuchi, 2011), hardware implementations of neural systems suffer from the same problems related to parameter and state estimation. It is very difficult to estimate the parameters of hardware neural systems due to their nonlinear relationships to the parameters of the model and their nonidealities caused by device mismatch, noise, and temperature sensitivity (Pavasović, Andreou, & Westgate, 1994). One fundamental difference between biological neural systems and hardware-engineered systems, however, is that the blueprint of the studied neuron is known by design. This strategically places hardware-engineered devices between software simulations and their biological counterparts. Neuromorphic systems are therefore ideal testing platforms for numerical methods of parameter estimation. Furthermore, DSPE techniques can be very useful for configuring neuromorphic systems to implement a desired set of behaviors (Camilleri, Giulioni, Mattia, Braun, & Del Giudice, 2010; Neftci, Chicca, Indiveri, & Douglas, 2011). In this work, we apply DSPE on data collected from neural populations of VLSI neurons.

The remainder of the letter is organized as follows. First, in section 2, we give an overview of the DSPE method. In section 3.1, we describe a control experiment where we validate the method on a fully known model using data generated from software simulations. Unlike for the neuromorphic VLSI system, the parameters of the simulated system are all perfectly known and therefore allow direct, quantitative comparisons with the results provided by DSPE. In section 3.2, we apply the method to the neuromorphic system and compare the results obtained with those of previous parameter estimation techniques applied to the same system. The findings are then discussed in section 4.

2 Materials and Methods

2.1 The Dynamic State and Parameter Estimation Method. The DSPE method (Creveling, Gill, & Abarbanel, 2008) is an optimization technique that uses principles from observer theory (Nijmeijer, 2001) and synchronization (Pecora & Carroll, 1990) of nonlinear systems to determine parameters and unmeasured state variables in an experimental system. DSPE allows one to derive the properties of a nonlinear dynamical system using only experimental measurements coupled with a mathematical model of the studied system. Specifically, DSPE is implemented by synchronizing the output of experimental observations of a system with a model of the system. This is a balanced synchronization between the data and the model, in the sense that the experimental data are transmitted to the model accurately and efficiently, yet the coupling is not so strong that the model becomes overwhelmed by the data. Twin experiments, where the data have been created using a computational model, have shown that the technique works reliably for spiking neuron models (e.g., Hodgkin-Huxley), as well as nonlinear circuits, and simple geophysical fluid flow models.

The formulation of DSPE is as follows. The state of an experimental system is described by N independent dynamical variables, $\mathbf{x}(t) = [x_1(t), x_2(t), \ldots, x_N(t)]$, $\mathbf{P}$ fixed parameters, and eventually external inputs to the system. To determine the full state of the system, the $\mathbf{x}(t)$ are all required, but typically only one or a few components can be observed. If L state variables can be observed, $\mathbf{x}(t) = [x_1(t), x_2(t), \ldots, x_L(t), \mathbf{x}_\perp(t)]$, where $\mathbf{x}_\perp(t)$ are unmeasured state variables, and the first-order differential equations describing the model are

$$\frac{dx_1(t)}{dt} = G_1(x_1(t), x_2(t), \ldots, x_L(t), \mathbf{x}_\perp(t), \mathbf{P}),$$

$$\frac{dx_2(t)}{dt} = G_2(x_1(t), x_2(t), \ldots, x_L(t), \mathbf{x}_\perp(t), \mathbf{P}),$$

$$\vdots$$

$$\frac{\mathrm{d}x_L(t)}{\mathrm{d}t} = G_L(x_1(t), x_2(t), \ldots, x_L(t), \mathbf{x}_\perp(t), \mathbf{P}),$$
$$\frac{\mathrm{d}\mathbf{x}_\perp(t)}{\mathrm{d}t} = \mathbf{G}_\perp(x_1(t), x_2(t), \ldots, x_L(t), \mathbf{x}_\perp(t), \mathbf{P}). \tag{2.1}$$

A typical solution to this type of problem is a least-squares optimization of the error between the measured data, $x_1(t)$ and a model $y_1(t)$; this method works well for linear systems but breaks down for nonlinear systems with positive conditional Lyapunov exponents (CLEs) (Pecora & Carroll, 1990; Abarbanel et al., 2009). To combat this problem, the experimental data are coupled to a model system, $\mathbf{y}(t)$, with parameters $\mathbf{Q}$, as for an optimal-tracking problem. These coupling terms $\mathbf{u} = [u_1(t), u_2(t), \ldots, u_N(t)]$ drive the model system to synchronize with the data and reduce the CLEs to nonpositive values.

$$\frac{\mathrm{d}y_1(t)}{\mathrm{d}t} = F_1(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_1(t)(x_1(t) - y_1(t)),$$
$$\frac{\mathrm{d}y_2(t)}{\mathrm{d}t} = F_2(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_2(t)(x_2(t) - y_2(t)),$$
$$\vdots$$
$$\frac{\mathrm{d}y_L(t)}{\mathrm{d}t} = F_L(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_L(t)(x_L(t) - y_L(t)),$$
$$\frac{\mathrm{d}\mathbf{y}_\perp(t)}{\mathrm{d}t} = \mathbf{F}_\perp(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}). \tag{2.2}$$

Again, least-squares optimization could be used on the error between the model and the measured data over the course of the time series, but the addition of the coupling terms, $\mathbf{u}(t)$, complicates matters. This coupling must be chosen large enough to cause synchronization of the data to the model (and eliminate the positive CLEs) but must not overwhelm the underlying dynamics of the system. The addition of the coupling term into the cost function for the optimization will ensure that the coupling does not become too large, while appropriate bounds for the range of the coupling ensure that it becomes large enough. Therefore, the optimization to be performed involves the minimization of

$$C(\mathbf{y}, \mathbf{u}, \mathbf{Q}) = \frac{1}{2T}\int_0^T \sum_{i=1}^{L}\{(x_i(t) - y_i(t))^2 + u_i(t)^2\}\mathrm{d}t,$$

subject to:

$$
\begin{aligned}
0 &= -\frac{dy_1(t)}{dt} + F_1(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_1(t)(x_1(t) - y_1(t)),\\
0 &= -\frac{dy_2(t)}{dt} + F_2(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_2(t)(x_2(t) - y_2(t)),\\
&\vdots\\
0 &= -\frac{dy_L(t)}{dt} + F_L(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}) + u_L(t)(x_L(t) - y_L(t)),\\
0 &= -\frac{d\mathbf{y}_\perp(t)}{dt} + \mathbf{F}_\perp(y_1(t), y_2(t), \ldots, y_L(t), \mathbf{y}_\perp(t), \mathbf{Q}),
\end{aligned}
\tag{2.3}
$$

and also subject to suitable bounds for the state variables and parameters.

Since the experimental data are discrete, the cost function integral becomes a summation, and the constraint equations are transformed to a discrete integration map using a numerical integration algorithm (e.g., Simpson's rule), with each time step having its own constraint equation for each state variable. The unknown variables for the optimization are each of the state variables and the couplings at each time step and each of the parameters.

A variety of optimization software and algorithms is available to solve this problem. IPOPT (Wächter & Biegler, 2006) was chosen since it is widely available, is designed for nonlinear problems with sparse Jacobian structure, can be parallelized with the linear solver PARDISO (Schenk & Gärtner, 2004), and can thus solve large problems in a reasonable time. Depending on the problem and the data set, a few thousand data points are necessary to explore the state space of the model and allow DSPE to produce accurate solutions. For problems on the VLSI neural networks described here, which involve the estimation of six parameters and approximately 1000 data points, this results in tens of thousands of constraints and unknown variables, which can be cumbersome for some solvers. Due to the discretized structure of the problem, however, the Jacobian of the constraints is sparse, and IPOPT can take advantage of this sparsity. The DSPE method can be used to determine the parameters of a neural network and infer the parameters of the single neurons composing this network. In the next section, we show that this method is also efficient for determining the parameters of a VLSI network comprising several integrate-and-fire (I&F) neurons.

2.2 Experimental Setup: The VLSI Neuromorphic System. Recent implementations of neuromorphic multineuron chips comprise large assemblies of biologically plausible silicon neurons and synapses with biophysically realistic dynamics (Indiveri et al., 2011). These devices permit the implementation of large, massively parallel networks in VLSI. We next

present the neuron model used in the neuromorphic VLSI chip and then give a description of the network of VLSI neurons and its dynamics.

2.2.1 The Single VLSI I&F Neuron. The VLSI neuron circuit considered here implements an I&F model featuring positive feedback, constant leakage, refractory period, and thousands of plastic and nonplastic synapses. It has been extensively described and characterized in Indiveri et al. (2006), Ben Dayan Rubin, Chicca, and Indiveri (2004), and Neftci et al. (2011). Here, we summarize its relevant properties. The dynamics governing the membrane potential V_m below firing threshold obey the differential equation

$$C\frac{\mathrm{d}}{\mathrm{d}t}V_m = I(t) - \beta + I_{fb}\exp\left(\frac{V_m - V_{th}}{\Delta_T}\right), \tag{2.4}$$

where C is the membrane capacitance, $I(t)$ an input current, β represents a constant leak, I_{fb} the positive feedback current factor, V_{th} the positive feedback activation voltage, and Δ_T the slope factor that determines the sharpness of the spike. When the membrane potential reaches its firing threshold an address-event is produced, and a digital event is transmitted off-chip using the address event representation (AER) communication protocol (Lazzaro, Wawrzynek, Mahowald, Sivilotti, & Gillespie, 1993). A resetting circuit then pulls the membrane potential V_m down to the reset potential and clamps it there during a refractory period τ_{ref}.

2.2.2 Networks of VLSI I&F Neurons. The VLSI neurons in our chip are connected to each other through synapses with biophysically realistic dynamics. Our goal is to derive a model that accurately describes this spiking neural network. Unfortunately, the firing and reset mechanism used in the model above yields a system that does not have continuous dynamics, a requirement of the proposed DSPE method. One possibility is to model the (continuous time) detailed transistor equations of the circuit. However, this would introduce interactions at very different timescales and render the problem intractable when several interacting neurons are considered.

To circumvent this problem, the VLSI spiking neural network can be modeled using a mean-field approach. In this approach, the state variables represent collective dynamics of the neurons composing the population via analytically derived or experimentally estimated activation functions (Wilson & Cowan, 1972; Abbott, 1994; Ben-Yishai, Lev Bar-Or, & Sompolinsky, 1995; Yuille & Grzywacz, 1989; Fusi and Mattia, 1999). When applied to collections of I&F neurons, the mean-field approach provides a continuous system. The stochastic nature of the inputs and the variability among the VLSI neurons (e.g., due to device mismatch) cause the resulting system to be stochastic. Because such systems are too difficult to solve in the general case, the collective dynamics of the system is often studied in the

diffusion approximation (Wang, 1999; Brunel & Hakim, 1999; Brunel, 2000b; Fusi and Mattia, 1999; Renart, Brunel, & Wang, 2003; Deco, Jirsa, Robinson, Breakspear, & Friston, 2008). In this approximation, the firing rates of individual neurons are replaced by a common time-dependent population activity variable with the same mean and two-point correlation function as the original variables. The approximation is appropriate when the effects of the synaptic inputs are very small but the overall firing rate is very high. For our experimental system, this means the charge delivered by each spike to the postsynaptic neuron is small compared to the typical charge necessary to make the neuron spike, there is a large number of afferent inputs to each neuron, and the spike times are uncorrelated. The parameters of the VLSI I&F neurons can be configured to verify the above requirements, so we can use a mean-field model to represent the average state of a collection of VLSI neurons. For the networks presented in this letter, we have $q \leq 0.3C\Theta$, where q is the synaptic weight, $\Theta \cong 1.2$ V is the approximate firing threshold of the neurons (Neftci et al., 2011), and $C = 1$ pF is the membrane capacitance, with around 20% connectivity.

2.2.3 Neural States and Dynamics. In the neuromorphic system used for this work, the slowest dynamics are usually those of the synapses. Therefore, we describe a model in which the synaptic states have temporal dynamics and the other quantities are considered instantaneous. The synaptic currents depend on the activities of the presynaptic neurons and are commonly modeled using first-order linear filter dynamics (Destexhe, Mainen, & Sejnowski, 1998). The average postsynaptic current (PSC) of a synapse of weight q and time constant τ in response to an arbitrary spike train of firing rate ν can be modeled as

$$\tau_{ij}\frac{\mathrm{d}}{\mathrm{d}t}I_{syn_{ij}}(t) + I_{syn_{ij}}(t) \cong q_{ij}\nu_i(t), \tag{2.5}$$

where ν_i refers to the firing rate of neuron i. If the time constant of the membrane is much smaller than the time constant of the synapse (as is the case for the VLSI neurons used in this work), the firing rate ν_i in response to both current injection I_{inj_i} and a synaptic input $\sum_j I_{syn_{ij}}$ is approximately equal to

$$\nu_i(t) \cong \sigma\left(\sum_j I_{syn_{ij}} + I_{inj_i} - \beta\right), \tag{2.6}$$

where $\sigma(\cdot)$ is the activation function of the neuron. The activation function is threshold linear and saturates at high values due to an absolute refractory

period τ_{ref}, during which the neuron does not fire. In the diffusion approximation, the form of σ can be analytically derived (Fusi & Mattia, 1999) and its parameters fit by experimental measurements. In the regime of interest, where the synaptic time constants are long compared to the typical interspike intervals of the active neurons (5–50 ms), the currents impinging on the neurons have small variances compared to their mean, so we can focus on the signal-driven regime. Here, we use a slight approximation of Fusi and Mattia's analytically derived activation function in this regime (as opposed to a noise-driven one), which is more stable in numerical optimizations:

$$[x]^+ = x + \alpha^{-1}\log(1 + e^{-\alpha x}),$$
$$\sigma(x) = \frac{[wx - T]^+}{1 + \tau_{ref}[wx - T]^+}, \tag{2.7}$$

where w, T, τ_{ref}, and α are parameters adjusted to the neural populations prior to running the DSPE. The parameter w is the weight of the external AER synapse, T is proportional to the somatic leak of the neurons (β in equation 2.4), τ_{ref} is the refractory period of the neurons, and α sets the smoothness of the threshold. The threshold-linear function $[\cdot]^+$ is equal to the integral of the sigmoid function (Chen & Mangasarian, 1995). It behaves well in numerical methods because its derivative lies between 0 and 1. In addition, the smooth onset of the threshold function approximates the effect of the noise at low frequencies.

The state of the synapse can be described by a single, dimensionless synaptic variable s_i, where $I_{syn_i} = s_i \frac{q_i}{\tau_i}$. Combining equations 2.5 and 2.6 results in a closed-form solution consisting of N^2 equations. This solution can be simplified to N equations if the synapses of each neuron have the same dynamics (i.e., when all the time constants of efferent synapses are equal: $\tau_{ij} = \tau_i$, $\forall j$),

$$\frac{\mathrm{d}}{\mathrm{d}t}s_i + \frac{s_i}{\tau_i} = \sigma\left(I_{inj,i} + \sum_j \frac{q_{ij}}{\tau_j}s_j - \beta_i\right) \quad \forall i = 1, \ldots, N. \tag{2.8}$$

These N equations can describe the firing rate for any network of neurons with the dynamics provided above and with long synaptic time constants compared to the neural interspike intervals.

2.2.4 An Excitatory-Inhibitory Network. The simplest model able to exhibit the dynamics of a generic multineuron system composed of excitatory and inhibitory neurons is the excitatory-inhibitory network (EIN; Dayan & Abbott, 2001). The EIN consists of two coupled populations of spiking

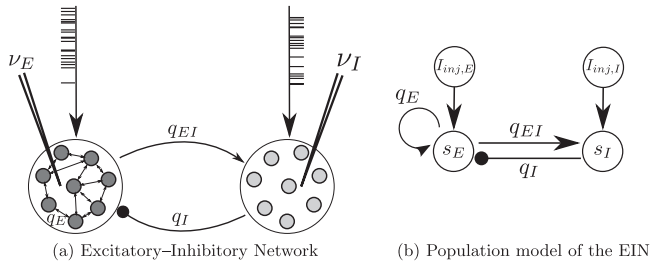

(a) Excitatory–Inhibitory Network (b) Population model of the EIN

Figure 1: Model of the EIN of spiking neurons. (a) The neurons in the excitatory population (left circle) have local, nearest-neighbor recurrent couplings among each other and excite the neurons within the inhibitory population (right circle), which in turn inhibit the excitatory neurons. (b) The EIN can be modeled using two recurrently coupled units whose variables represent the synaptic variable. The mean firing rate of the neurons can be inferred from the synaptic variables using equation 2.6.

neurons: one population is excitatory and the other is inhibitory. The neurons in the excitatory population have recurrent couplings among each other and excite the neurons within the inhibitory population, which in turn inhibit the excitatory neurons. The recurrent excitatory couplings were chosen to reflect those of the VLSI chip used in this work and consisted of first, second, and third nearest-neighbor connections on a 1D topology (see Figure 1a). These excitatory and inhibitory couplings were originally designed to generate soft winner-take-all dynamics (Yuille & Geiger, 2003).

We can use equation 2.8 to model this network. However, this results in N differential equations (one for each type of synapse), which can be cumbersome for the DSPE when N is large. Instead, we use a mean-field approximation in which we replace the state variables of a collection of identical targets receiving statistically identical inputs by a single variable. This approximation is accurate if the following two conditions are true: the external inputs are chosen to be independent and Poisson distributed, with identical mean rates for every neuron in a common population, and the conditions described for the diffusion approximation are held. In this case, the state of the synapses in a common population can be described by a single differential equation. This means that the EIN dynamics can be expressed with only two variables, s_E and s_I, representing the collective dynamics of the excitatory synapses and the inhibitory synapses, respectively. The modeled system is illustrated in Figure 1b. To take into account a nonlinearity present in the VLSI implementation of the inhibitory synapse, we introduce a threshold-linear term, $\sigma_{inh}(x - T_{Isyn}) \cong [x - T_{Isyn}]^+$ (Bartolozzi & Indiveri, 2007), where T_{Isyn} is the free parameter of this corrective term and $[\cdot]^+$ is a half-wave rectification (see equation 2.7). With this nonlinearity, the

dynamics governing s_E and s_I become

$$\begin{aligned}
&\frac{\mathrm{d}}{\mathrm{d}t}s_E + \frac{s_E}{\tau_E} = \nu_E = \sigma\left(I_{inj,E} - \beta_E + \frac{\bar{q}_E}{\tau_E}s_E - \frac{\bar{q}_I}{\tau_I}\sigma_{inh}(s_I - T_{Isyn})\right),\\
&\frac{\mathrm{d}}{\mathrm{d}t}s_I + \frac{s_I}{\tau_I} = \nu_I = \sigma\left(I_{inj,I} - \beta_I + \frac{\bar{q}_{EI}}{\tau_k}s_E\right),\\
&\text{with} \quad \bar{q}_k = N_{q_k}q_k, \;\; k = E, EI, IE,
\end{aligned} \tag{2.9}$$

where $\bar{q}_E$, $\bar{q}_{EI}$, and $\bar{q}_I$ are the mean recurrent excitatory and inhibitory synapse weights in the network, the N_{qk}'s are the number of synapses of type k per neuron (for the neuromorphic system considered here, $N_{qE} = 6$ corresponding to the number of nearest neighbor synapses, $N_{qEI} = 30$ and $N_{qIE} = 4$), τ_E, τ_I are the synaptic time constants, and β_E, β_I represent the leak in the somata. The neurons in the chip communicate with the outside world via AER events. AER synapses can be stimulated externally using a PC or another VLSI chip. The currents delivered by these external synapses are represented here by the injection currents $I_{inj,E}$ and $I_{inj,I}$.

The interaction between the two populations in the EIN gives rise to highly nonlinear dynamics, which is fit using DSPE. Because the synaptic variables s_E and s_I cannot be measured directly, the implementation of our DSPE method includes dynamics for the "readout variables" ν_E and ν_I, which can be directly measured by average population rates.

A similar mean-field approach was used to model hardware spiking neural networks by Camilleri et al. (2010).

2.3 Dynamic State and Parameter Estimation Experimental Procedure. In order to estimate the parameters of the system, we must define a procedure that explores the state space of the system. The DSPE procedure can be subdivided into three components: (1) an experiment that explores the state space of the experimental system in terms of (2) a model that can be described in terms of coupled differential equations and (3) an implementation of the problem to be solved using an optimizing program.

The procedure for implementing DSPE of the EIN is the following. First, Each population of the EIN is stimulated with inhomogeneous Poisson spike trains whose mean rates are dictated by a common Ornstein-Uhlenbeck (OU) process driven by white noise, $\tau\frac{\mathrm{d}}{\mathrm{d}t}b = -b + A + \sigma\xi$, where ξ is gaussian white noise. Firing rates can only be positive, so we take $\max(b, 0)$ as the rate for the Poisson process. The values of the time constant τ, the base rate A, and the amplitude σ are specific to each experiment and are indicated in the relevant sections. Second, the input and output mean activity of each population is computed by convolving each spike with an exponentially decaying kernel of time constant 50 ms and taking

the population average. The obtained data are sampled at 200 Hz. Finally, the model is transformed into a DSPE instance with the use of Python scripts constructed for this purpose. The Python scripts take the continuous vector field, discretize according to Simpson's integration rule, and output the discretized vector field along with its Jacobian (first derivatives) and Hessian (second derivatives) matrices to a collection of C++ programs that call the appropriate IPOPT libraries to perform the DSPE optimization (Toth, 2011). The EIN system was implemented with the control terms acting on the firing rate variables ν_E and ν_I.

In terms of equations 2.2, the observed states are given by $x_1(t) = \bar{\nu}_E(t)$, $x_2(t) = \bar{\nu}_I(t)$, where the bar denotes experimentally measured data. The control terms are $u_1(t)$ and $u_2(t)$, respectively for ν_E and ν_I. The state of the system is described by four dimensions: $\mathbf{y}(t) = [s_E(t), s_I(t), \nu_E(t), \nu_I(t)]$. We assumed that the synaptic dynamics were the slowest in the overall system, such that the dynamics of ν could be considered instantaneous. Under these assumptions, the differential equations implementing the DSPE problem become

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}s_E &= -\frac{s_E}{\tau_E} + \nu_E,\\
\frac{\mathrm{d}}{\mathrm{d}t}s_I &= -\frac{s_I}{\tau_I} + \nu_I,\\
\nu_E &= \sigma\left(I_{inj,E} - \beta_E + \frac{\bar{q}_E}{\tau_E}s_E - \frac{\bar{q}_I}{\tau_I}\sigma_{inh}(s_I - T_{Isyn})\right) + u_1(t)(\bar{\nu}_E - \nu_E),\\
\nu_I &= \sigma\left(I_{inj,I} - \beta_I + \frac{\bar{q}_{EI}}{\tau_E}s_E\right) + u_2(t)(\bar{\nu}_I - \nu_I). \qquad (2.10)
\end{aligned}$$

For technical reasons, ν was implemented with temporal dynamics, but with a time constant much faster then the synaptic one (2 ms), such that it did not affect the final result. This differential system is provided to the DSPE, along with the boundaries on the values of the parameters and the states.[1]

3 Results

3.1 Validation: Software-Simulated Neurons. To validate the DSPE method proposed in this work, we can apply it to a system whose parameters are completely known beforehand and compare them with the fitted

[1]The configuration files used for DSPE are provided online at http://ncs.ethz.ch/internal/files/dpe/at_download/file or can be requested by e-mailing the authors.

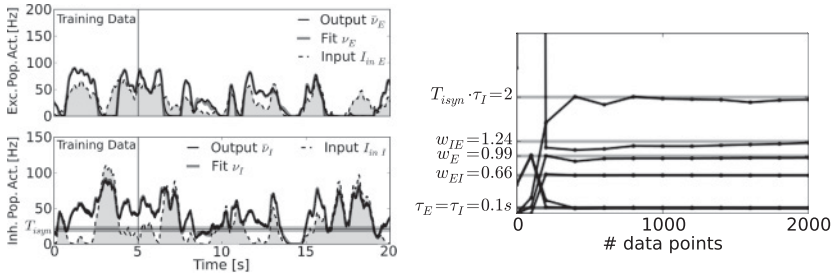

Figure 2: DSPE of a software-simulated EIN. The EIN simulation comprises one excitatory population of 30 neurons and one inhibitory population of 4 neurons. The input consists of 20 s of inhomogeneous Poisson spike trains. The signals governing their instantaneous firing rates are drawn from a common OU process driven by white noise (time constant $\tau = 500$ ms, amplitude $\sigma = 150$ Hz, plus a base firing rate of $A = 100$ Hz). (a) The resulting input firing rates (shaded curve) along with the observed and predicted EIN firing rates (black line and thick gray line, respectively). To allow easier comparison between input and output curves, the input firing rates presented here are the result of the multiplication of the measured input rate with the weight of the input synapses ($w_{AER,E} = 0.197$ and $w_{AER,I} = 0.214$). The training data used for the DSPE consist of 1000 data points, corresponding to 5 s of simulated data (vertical line). From the vertical line on, the thick line shows the prediction obtained using the fitted parameters. (b) Estimated parameters in dependence of the data set size. Each point indicates the value of one of the six parameters estimated by the DSPE. These curves shows that the fitted parameters did not change significantly beyond 400 data points. The given dimensionless weight parameters w_i, $i = E, EI, IE$ are related to the synaptic weights $\bar{q}_i$ by $w_i = \frac{\bar{q}_i}{C\Theta}$. The values of other fixed parameters used in the simulation were $C = 1$ pF, $\beta = 4.15$ pA.

parameters. For this, we carry out a software simulation of the EIN, implemented using the BRIAN Neural Simulator (Goodman & Brette, 2008). The simulated network consisted of I&F neurons models of the type described in equation 2.4 but without the feedback term. Instead, the neuron was considered to fire when it reached $\Theta = 1.2$ V. The other unfitted parameters of the simulated I&F neurons (such as capacitance and refractory periods) were set to values comparable to those of the VLSI I&F neurons. The excitatory population and the inhibitory population consisted of 30 and 4 I&F neurons, respectively. These numbers were chosen to match those of the hardware experimental system (see section 3.2). Both populations were stimulated by Poisson spike trains whose mean rates followed an OU process (see section 2.3). For each population, the input, the measured output, and the predicted output are shown in Figure 2a.

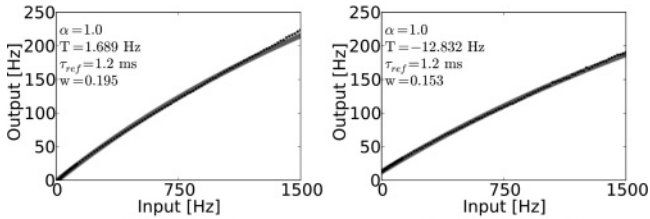

Figure 3: Experimentally measured activation functions. We estimate the activation function of each population of VLSI neurons by measuring their responses to constant mean rate regular spikes trains. The saturation at large output firing rates is due to the refractory period in the neurons. We fitted the measured input-output rates with the activation function provided in equation 2.7 (thick line).

We used 1000 data points to estimate the parameters of the model, corresponding to 5 s of simulated data (see Figure 2a, up to the vertical line). The fits matched the simulated data with nearly perfect accuracy (see Figure 2a, thick gray line versus black line). The data not used for the DSPE, from the vertical line on 5 to 20 s, were equally well predicted by the fitted model.

To certify that the dynamic repertoire used to determine the system's parameters was wide enough to account for all the observable dynamic regimes, we carried out the DSPE over different sizes of the same data set and compared the fitted parameters to the ones used to generate the data. The results are summarized in Figure 2b, where the points on each curve indicate the values of the estimated parameters for the given data set sizes. We see that 400 data points were sufficient for the DSPE algorithm to converge to the correct parameters (see Figure 2b, horizontal lines). The fit improved only very slightly by increasing the amount of data used for the estimation. The slight discrepancies between the true and estimated values of η_E and T_{Isyn} remained even for the largest data set we have considered. This was due to the mean-field approximation of the simulated neural system and the relatively small number of neurons, which made this approximation less accurate.

These results demonstrate that the considered experimental protocol is adequate for estimating the parameters of the chosen model. We can now address the problem of estimating hardware neuron parameters.

3.2 Real Experiment: VLSI Network Activity. To estimate the parameters of the VLSI EIN, we apply DSPE using a similar experimental protocol as described in the previous section. The activation functions of the excitatory and the inhibitory populations were estimated using steady-state input-output activity measurements (see Figure 3).

To show that the results of the DSPE generalize well, we have recorded the response of the VLSI EIN for several input conditions and separated them into two sets: the first set was used as training data for the DSPE, and the second set was used as a testing set to check the accuracy of the fitted model and was not used for the DSPE algorithm. The training set consisted of the same type of input as for the software-simulated EIN: an OU process driven by white noise with time constant $\tau = 500$ ms. The testing data consisted of three sets resulting from inhomogeneous Poisson inputs modulated by three different signals: one OU process (different realization than training data), one sinusoidal (0.76 Hz), and one consisting of box-shaped pulses of varying heights and durations.

We observe that the model fitted by the DSPE faithfully reproduced the training data (see Figure 4a). The fit with the testing sets were equally good (see Figure 4b–4d), but less so for high-frequency variations (e.g, see the activity peaks in Figures 4b and 4d). The possible explanation for this discrepancy is due to the inability of the mean-field model to describe all the complexity of the underlying VLSI neural system. In fact, the raster plot in Figure 5 shows that winning regions of activity emerged in the VLSI EIN. This effect was caused by the nearest-neighbor connections between the excitatory neurons and is not taken into account by our mean-field model.

The fitted model was able to predict the settlement to an attractor state, as observed in Figures 4b and 4c by the ongoing, persistent activity in the absence of inputs. The emergence of attractor states in spiking neural networks is a widely studied phenomenon (Amit & Brunel, 1997b; Brunel & Wang, 2001; Renart et al., 2003) and was previously reported in this chip (Neftci, Chicca, Cook, Indiveri, & Douglas, 2010), and other neuromorphic multineuron chips (Camilleri et al., 2010; Massoud & Horiuchi, 2011).

3.3 A Systematic Parameter Estimation Technique. Previous work with this type of hardware has described a parameter translation method in which the parameters of single neurons are mapped to those of the VLSI I&F. This required a calibration procedure based on population gain measurements to be carried out for each parameter (Neftci et al., 2011). The DSPE method described here has the advantage that it can estimate all the parameters of the system in one shot. We show here that the estimated parameters are comparable to those provided by the parameter translation method of Neftci et al. (2011). For this, we run the experiment described in section 3.2 for a large number of parameters sets by varying the excitatory couplings q_E while the other parameters were kept constant. The input consisted of Poisson spike trains whose rate followed an OU process ($\tau = 500$ ms, $\sigma = 150$ Hz, $A_E = 150$ Hz, $A_I = 90$ Hz) and was identical for each DSPE run. The results of both methods are shown in Figure 6 and Table 1. The line is the result obtained by direct parameter translations. Its slope represents the subthreshold slope factor of the synapse's weight transistor and the value of the line at 0 V is approximately the product

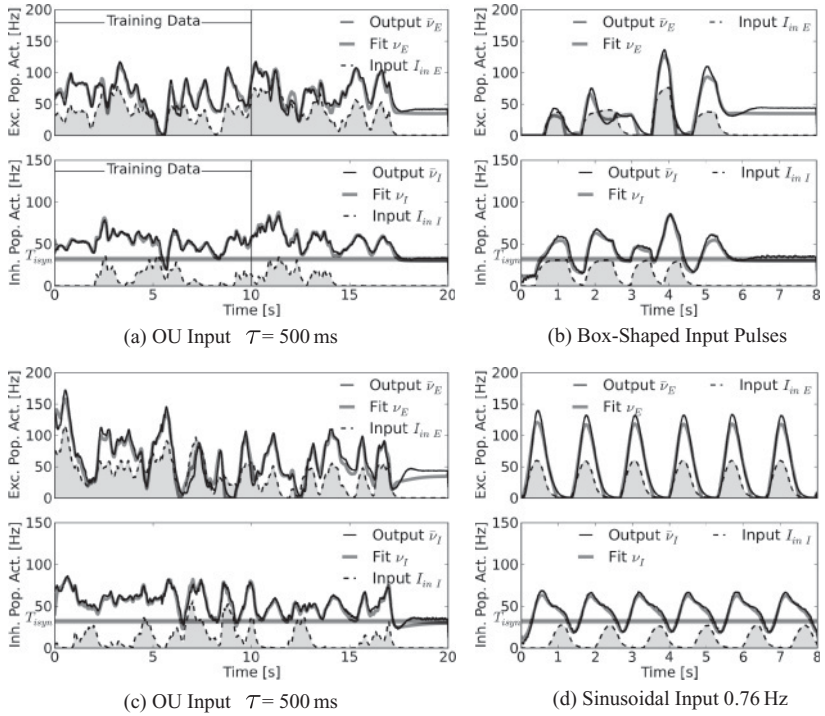

Figure 4: DSPE applied to the VLSI EIN. As with the software-simulated system, the input used to explore the system's dynamics consisted of inhomogeneous Poisson spike trains with mean rates governed by an OU process. The parameters of the process were set to $\tau = 500$ ms, $\sigma = 150$ Hz, and $A_E = 180$ Hz for the excitatory neurons and $A_I = 60$ Hz for the inhibitory neurons. For presentation purposes, the input firing rates were multiplied by the weight of the input synapses estimated by the transfer curves of Figure 3 ($w_{AER,E} = 0.195$ and $w_{AER,I} = 0.153$). The training data consisted of 10 s of recorded output data (up to the vertical line), sampled at 200 Hz. (a) Both the training data and the following 10 s of recorded data not used for the estimation were faithfully reproduced using the fitted model. (b–d) Comparison of the responses of the fitted model to those of the VLSI EIN for three different input types, which were not used for estimation. In panels a, b, and c, we observe the emergence of activity persisting even in the absence of input. This effect was well accounted by the fitted model, even though no persistent activity state was present in the training data set.

of its off-current and a pulse duration resulting from an incoming spike (Bartolozzi & Indiveri, 2007). Although the latter cannot be measured in this chip, its rough estimation to a few μs leads to off-current values compatible with those measured by the method described in Neftci et al. (2011).

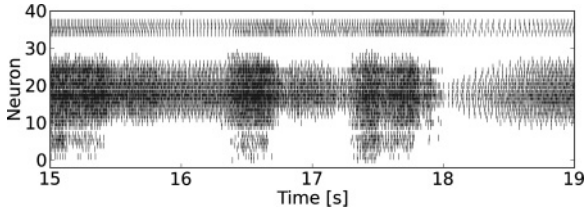

Figure 5: Example raster plot of the VLSI EIN. This raster plot shows an extract of the spiking activity of the data presented in Figure 4c, from time 15 s to 20 s. Neurons 1 to 30 are excitatory neurons, and neurons 34 to 38 are the inhibitory neurons. Due to the recurrent couplings in the EIN, a group of winning neurons emerged around position 20. The exact position of the winner is mainly determined by the neural inhomogeneities in the network (due to fabrication mismatch), visibly favoring the group of neurons around 20 over the others. Such effects are not taken into account by our mean-field model and is most likely at the origin of the discrepancies observed in Figure 4.

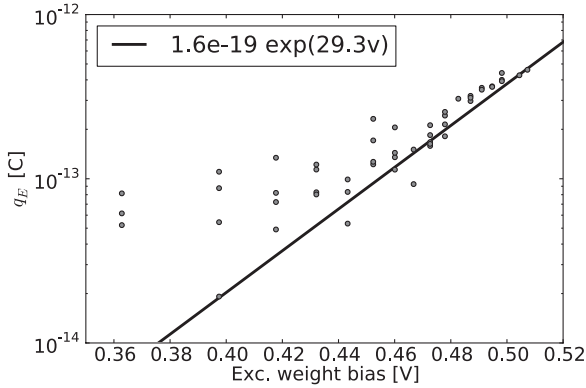

Figure 6: Systematic parameter estimation of the excitatory couplings and comparison to the parameter translation technique. The experiment described in section 3.2 is run for a range of excitatory weight biases controlling the parameter q_E, while the other biases were kept constant. During the bias sweep, the parameters τ_I, τ_I, T_{Isyn}, q_{IE}, and q_{EI} were kept constant but were nevertheless fitted during each run of the DSPE. The fitted parameters were accepted when the DSPE had converged (approximately 75% of the cases), and the final value of the cost function was below an acceptable threshold. They are compared to the results of the parameter translation method presented in Neftci et al. (2011) (line). For a comparison of the unvaried parameters, see Table 1.

Table 1: Comparison Between Parameters Obtained from Neftci et al. (2011) and the Parameters Estimated from the VLSI I&F Neurons Using DSPE.

Parameter	Parameter Translations	DSPE
$\bar{q}_{EI}$	3.89×10^{-14} C	$4.265 \pm 0.495 \times 10^{-14}$ C
$\bar{q}_{IE}$	0.8×10^{-13} C	$2.000 \pm 0.808 \times 10^{-13}$ C
τ_I	0.111 s	0.0512 ± 0.0120 s
τ_E	0.027 s	0.0304 ± 0.0043 s
T_{Isyn}	$\cong 4.8 \times 10^{-11}$ A	$5.068 \pm 1.848 \times 10^{-11}$ A

The matching between both methods was very good when the weights were significant ($\bar{q}_E > 1 \cdot 10^{-13}$ C, corresponding roughly to a synaptic efficacy higher than 0.1, e.g., 10 inputs spikes generates 1 output spike). A clear discrepancy is observable for the inhibitory synapse time constant τ_I and weight q_{IE} (see Table 1). We believe that the DSPE method described here produces more accurate estimations than the direct parameter translation method of Neftci et al. (2011) because it is based on the direct measurement of the system response.

4 Discussion

Neuroscientists often propose detailed biophysical models to probe the computational purposes of their studied neural systems. To determine the parameters of their model, electrophysiologists commonly employ heuristic techniques such as current pulses, current ramps, and conductance injections. Unfortunately, the estimation of the neural properties of biological neurons from experimental data becomes very difficult when the model used involves many parameters and states that cannot be simultaneously measured. Furthermore, due to simplifying assumptions it is not possible to generalize much of the existing work on parameter estimation (Brillinger, 1988; Huys, Ahrens, & Paninski, 2006; Paninski et al., 2004; Okatan et al., 2005; Rossant et al., 2010) to large neural systems consisting of recurrently interacting neurons.

In this letter, we presented a method for nonlinear neural model validation and parameter estimation using a recently proposed DSPE technique by Abarbanel et al. (2009). We adapted the DSPE technique for the purpose of estimating the parameters of engineered neuromorphic VLSI systems that emulate the computational processes observed in biological neurons. The technique is applied to a VLSI neural network composed of 30 excitatory and 4 inhibitory spiking neurons and was able to accurately determine the network parameters without the use of membrane potential traces.

We have demonstrated the proposed method in three experiments. First, we validated the method on a fully known model using data generated from software simulations of the EIN network and observed that the inferred parameters were accurate. Then we determined the parameters of the VLSI system. The fitted model accurately reproduced both the training data and a testing data set, which we did not use during the DSPE optimization procedure. Finally, in order to show that DSPE can be used for a large number of parameter configurations, we estimated the parameters for a large range of parameter sets. This has revealed an exponential relationship between bias settings and fitted values, in agreement with the design of the chip and previous observations (Neftci et al., 2011).

Most investigators in the field of neuromorphic engineering employ manual or ad hoc methods to configure their devices. Manual calibration becomes intractable as the number of neurons in a single chip increases and some variables become impossible to observe directly (e.g., because of the limited number of pins on a chip). When the neurons in the population are all identical, we can exploit the relationship between the population models and the spiking neuron to infer mean, single-neuron parameters such as synaptic time constants and weights. Therefore, our choice for using a neuromorphic VLSI system was not only due to its practicality compared to biological systems, but also because parameter estimation methods are of great interest for the systematic configuration of neuromorphic hardware (Russell, Orchard, & Etienne-Cummings, 2007; Buhry, Saighi, Giremus, Grivel, & Renaud, 2009; Brüderle et al., 2009; Neftci et al., 2011; Camilleri et al., 2010). There are several cases where DSPE can be a very useful tool for neuromorphic systems. For example, Camilleri et al. (2010) implemented an attractor network model in VLSI spiking neural networks. To model the collective behavior of the neural populations, they used a mean-field approximation. The accuracy of their model was assessed by comparing the predicted transfer function to the one measured the VLSI neural network. The DSPE technique proposed here can be used as an additional tool to compare the dynamical properties of the neuromorphic system to those of the mean-field model. Another application is for the parameter translation method described in Neftci et al. (2011). This method requires a preliminary calibration procedure based on population gain measurements in steady state, carried out individually for each parameter to be estimated. The results of section 3.3 demonstrated that systematic parameter estimations were comparable to those of Neftci et al. (2011), suggesting that the proposed DSPE method can used for the calibration of the parameter translations.

The parameter translation method has an important difference compared to DSPE: the central assumption for configuring the VLSI system is that the effects of each coupling combine linearly. On the other hand, DSPE has the advantage that all the parameters can be simultaneously estimated a

posteriori. In this sense, it provides a measurement of the effective couplings in the system (Friston, 1994).

One requirement for applying DSPE is that the cost function, hence the dynamical system, must be continuous. This excludes the use of pure I&F neurons because of their hard reset operation. One possibility is to model the electronic circuit of the VLSI I&F neuron and the spiking mechanism explicitly, which is (at least) as difficult as modeling detailed sodium and potassium channel dynamics of H&H-like neuron models (Izhikevich, 2006). This approach would allow the consideration of neural networks whose coding strategy is based on spike timing and was previously achieved for single H&H neurons using DSPE in Abarbanel et al. (2009). However, it is very difficult to extend this result to networks of H&H and would likely require the simultaneous measurement of the membrane potentials of all the neurons in the network, which is often very difficult or even impossible in an experimental system. Instead, the presented estimation algorithm is based on a mean-field approach that provides a tractable description of the collective dynamics of I&F neurons. Mean-field methods can be used to model a large variety of neural systems. For instance, one can model conductance-based neurons (Shriki, Hansel, & Sompolinsky, 2003), adapting neurons, fast synapses (Fourcaud & Brunel, 2002), depressing and facilitating synapses (Barak & Tsodyks, 2007) and synaptic delays (Brunel, 2000a). Also, they can be formulated to take into account the heterogeneity of the neurons (Amit & Brunel, 1997a). For the purpose of determining synaptic weights, time constants, and thresholds, our experimental system is well approximated by the mean-field model consisting of two populations, one excitatory and one inhibitory (see equation 2.10). In our model, the dynamics are governed by the synapses when the synaptic time constant is large compared to the membrane time constant. As a result, the spiking statistics were quite regular and were driven mainly by the deterministic component of the input (Fusi & Mattia, 1999).

To apply the methods discussed in this letter, the assumptions of the mean-field model must be verified beforehand. The main disadvantage of this approach is due to the existence of many regimes in which these assumptions do not hold. In this case, the mean-field analysis becomes considerably more involved. The limitations of this approach were observable by the fact that the DSPE had failed to converge when the firing rate was too low because the activity was essentially driven by noise (Fusi & Mattia, 1999), the nearest-neighbor couplings broke the symmetry in the network (because it neglects the spatial correlations in the connectivity), or the state of the VLSI neural network went beyond its normal working range (e.g., due to runaway activity or transistors no longer operating in their intended regime). Some of these problems can be solved at the price of a more complex model. For instance, the first problem can be solved by taking the effect of noise into account in the activation function (Amit & Brunel, 1997a; Fusi & Mattia, 1999; Renart et al., 2003). The second problem can be solved by

considering a mean-field model with multiple excitatory populations, at the cost of an increase in the number of states.

The importance of determining confidence intervals of the estimated parameters is not to be underestimated. It is crucial for determining the accuracy of the tested model and is informative about the inevitable lack of sufficient state-space exploration, which can lead to poor generalization. For example, in the EIN studied in this letter, the inhibition has no effect until a certain threshold is reached. If the threshold is never reached (because of insufficient state-space exploration), the weight of the inhibition will be arbitrary because it will have no effect on the dynamics. This will lead to incorrect assumptions on the observed system. Among the previous work for parameter estimation, only those using maximum likelihood estimators such as Paninski et al. (2004), Huys et al. (2006), and Okatan et al. (2005) have statistical foundations and the confidence intervals can be obtained by calculating the variance of the estimator "for free" (Bishop, 2006). For all other parameter estimation techniques that are not built on statistical grounds (this includes the method presented in this letter), the confidence intervals must be estimated "brute force" by running the parameter estimation on several data sets or by using cross-validation techniques. This comes at the price of more experimental data and computational power.

The applications of the DSPE method presented in this letter are not limited to determining the parameters of engineered VLSI neural networks. When studied in conjunction with experimental data, the proposed DSPE will be of great value to validate a proposed model for neural masses and determine their interconnectivity (Freeman, 1975).

Another potential application of DSPE in the field of neural imaging is dynamic causal modeling (Friston, Harrison, & Penny, 2003), an increasingly popular technique for determining effective connectivity in neural systems, typically using fMRI or EEG signals (Stephan et al., 2007). The analysis of such data often requires the modeling of the underlying neural mechanism that generated the recorded signals (e.g., the balloon model; Buxton, Wong, & Frank, 1998). DSPE of neural masses can be used as the underlying engine for determining the parameters and the unmeasured variables of such experimental systems.

Our results in applying the DSPE procedure to the neuromorphic systems are very promising and suggest that DSPE may become a valuable tool for multineuron chip configuration methods. In future work, we will incorporate the DSPE procedure in an existing software framework developed for configuring the neural hardware (Sheik, Stefanini, Neftci, Chicca, & Indiveri, 2011).

Acknowledgments

This work was partially supported by the EU ICT Grant ICT-231168-SCANDLE (acoustic SCene ANalysis for Detecting Living Entities). This

work was partially supported by the National Science Foundation–sponsored Center for Theoretical Biological Physics at the University of California, San Diego. The research on numerical optimization was partially funded by NSF grant DMS-0511766. The chip was designed in collaboration with Elisabetta Chicca. We also thank the Neuromorphic Cognitive Systems group for their help in building the neuromorphic setup, Daniel Fasnacht for the design of the AER mapper board and the AER monitor and sequencer board, Rodney Douglas for discussion and support, Shih-Chi Liu for discussion and review, and the reviewers for their useful comments.

References

Abarbanel, H., Creveling, D., Farsian, R., & Kostuk, M. (2009). Dynamical state and parameter estimation. *SIAM Journal on Applied Dynamical Systems*, *8*, 1341–1381.

Abbott, L. (1994). Decoding neuronal firing and modeling neural networks. *Quarterly Review of Biophysics*, *27*, 291–331.

Amit, D., & Brunel, N. (1997a). Dynamics of a recurrent network of spiking neurons before and following learning. *Network: Computation in Neural Systems*, *8*(4), 373–404.

Amit, D., & Brunel, N. (1997b). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, *7*, 237–252.

Barak, O., & Tsodyks, M. (2007). Persistent activity in neural networks with dynamic synapses. *PLoS Computational Biology*, *3*(2), 323–332.

Bartolozzi, C., & Indiveri, G. (2007). Synaptic dynamics in analog VLSI. *Neural Computation, 19*(10), 2581–2603. http://ncs.ethz.ch/pubs/pdf/Bartolozzi_Indiveri07b.pdf.

Ben Dayan Rubin, D., Chicca, E., & Indiveri, G. (2004). Characterizing the firing proprieties of an adaptive analog VLSI neuron. In M. M. Auke Jan Ijspeert & N. Wakamiya (Eds.), *Biologically Inspired Approaches to Advanced Information Technology First International Workshop, Bioadit 2004*. New York: Springer-Verlag. http://ncs.ethz.ch/pubs/pdf/Ben-Dayan-Rubin_etal04.pdf.

Ben-Yishai, R., Lev Bar-Or, R., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences of the USA, 92*(9), pp. 3844–3848.

Bishop, C. (2006). *Pattern recognition and machine learning*. New York: Springer-Verlag.

Brette, R., & Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, *94*, 3637–3642.

Brillinger, D. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, *59*(3), 189–200.

Brüderle, D., Müller, E., Davison, A., Muller, E., Schemmel, J., & Meier, K. (2009). Establishing a novel modeling tool: A Python-based interface for a neuromorphic hardware system. *Frontiers in Neuroinformatics, 4*. http://frontiersin.org/Journal/Abstract.aspx?s=752&name=neuroinformatics&ART_DOI=10.3389/neuro.11.017.2009.

Brunel, N. (2000a). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, *8*(3), 183–208.

Brunel, N. (2000b). Persistent activity and the single-cell frequency–current curve in a cortical network model. *Network: Computation in Neural Systems*, *11*(4), 261–280.

Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, *11*(7), 1621–1671.

Brunel, N., & Wang, X. J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, *11*, 63–85.

Buhry, L., Saighi, S., Giremus, A., Grivel, E., & Renaud, S. (2009). Automated tuning of analog neuromimetic integrated circuits. In *Biomedical Circuits and Systems Conference BIOCAS 2009* (pp. 13–16). Piscataway, NJ: IEEE.

Buxton, R., Wong, E., & Frank, L. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance in Medicine*, *39*(6), 855–864.

Camilleri, P., Giulioni, M., Mattia, M., Braun, J., & Del Giudice, P. (2010). Self-sustained activity in attractor networks using neuromorphic VLSI. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks, 2010* (pp. 1–6). Piscataway, NJ: IEEE.

Chen, C., & Mangasarian, O. (1995). Smoothing methods for convex inequalities and linear complementarity problems. *Mathematical Programming*, *71*(1), 51–69.

Creveling, D. R., Gill, P. E., & Abarbanel, H. D. (2008). State and parameter estimation in nonlinear systems as an optimal tracking problem. *Physics Letters A*, *372*(15), 2640–2644. http://www.sciencedirect.com/science/article/B6TVM-4RH94RX-2/2/e087d04483f706ca711aa3c9aba883f7.

David, O., & Friston, K. (2003). A neural mass model for MEG/EEG: Coupling and neuronal dynamics. *NeuroImage*, *20*(3), 1743–1755.

Dayan, P., & Abbott, L. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.

Deco, G., Jirsa, V., Robinson, P., Breakspear, M., & Friston, K. (2008). The dynamic brain: From spiking neurons to neural masses and cortical fields. *PLoS Computional Biology*, *4*(8), e1000092.

Destexhe, A., Mainen, Z., & Sejnowski, T. (1998). Methods in neuronal modelling. In C. Koch & I. Segev (Eds.), *Methods in neuronal modeling* (pp. 1–25). Cambridge, MA: MIT Press.

Douglas, R., Mahowald, M., & Mead, C. (1995). Neuromorphic analogue VLSI. *Annu. Rev. Neurosci. 18*, 255–281.

Faugeras, O., Grimbert, F., & Slotine, J. (2008). Absolute stability and complete synchronization in a class of neural fields models. *SIAM Journal of Applied Mathematics*, *61*(1), 205–250.

Fourcaud, N., & Brunel, N. (2002). Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Computation*, *14*(9), 2057–2110.

Freeman, W. (1975). *Mass action in the nervous system*. New York: Academic Press.

Freeman, W., Ahlfors, S., & Menon, V. (2009). Combining fMRI with EEG and MEG in order to relate patterns of brain activity to cognition. *International Journal of Psychophysiology*, *73*(1), 43–52.

Friston, K. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, *2*(1–2), 56–78.

Friston, K., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, *19*(4), 1273–1302.

Fusi, S., & Mattia, M. (1999). Collective behavior of networks with linear (VLSI) integrate and fire neurons. *Neural Computation*, *11*, 633–652.

Goodman, D., & Brette, R. (2008). Brian: A simulator for spiking neural networks in Python. *Frontiers in Neuroinformatics*, *2*.

Hodgkin, A., & Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, *117*, 500–544.

Huys, Q., Ahrens, M., & Paninski, L. (2006). Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, *96*(2), 872.

Indiveri, G., Chicca, E., & Douglas, R. (2006). A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. *IEEE Transactions on Neural Networks*, *17*(1), 211–221. http://ncs.ethz.ch/pubs/pdf/Indiveri_etal06.pdf.

Indiveri, G., & Horiuchi, T. K. (2011). Frontiers in neuromorphic engineering. *Frontiers in Neuroscience*, *5*(0). http://www.frontiersin.org/Journal/FullText.aspx?s=755&name=neuromorphicengineering&ART_DOI=10.3389/fnins.2011.00118.

Indiveri, G., Linares-Barranco, B., Hamilton, T., Schaik, A. van, Etienne-Cummings, R., Delbruck, T., et al. (2011). Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience*, *5*, 1–23. http://www.frontiersin.org/Neuromorphic_Engineering/10.3389/fnins.2011.00073/abstract.

Izhikevich, E. (2006). *Dynamical systems in neuroscience: The geometry of excitability and bursting*. Cambridge, MA: MIT Press.

Jolivet, R., Lewis, T., & Gerstner, W. (2004). Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. *Journal of Neurophysiology*, *92*, 959–976.

Lazzaro, J., Wawrzynek, J., Mahowald, M., Sivilotti, M., & Gillespie, D. (1993). Silicon auditory processors as computer peripherals. *IEEE Transactions on Neural Networks*, *4*, 523–528.

Mahowald, M., & Douglas, R. (1991). A silicon neuron. *Nature*, *354*, 515–518.

Massoud, T., & Horiuchi, T. (2011). A neuromorphic VLSI head direction cell system. *IEEE Transactions on Circuits and Systems I*, *58*(1), 150–163.

Mead, C. (1989). *Analog VLSI and neural systems*. Reading, MA: Addison-Wesley.

Naud, R., Berger, T., Gerstner, W., Bathellier, B., & Carandini, M. (2009). Quantitative single-neuron modeling: Competition 2009. *Frontiers in Neuroinformatics*, 1–8. http://www.frontiersin.org/conferences/individual_abstract_listing.php?conferid=155&pap=2139&ind_abs=1&q=98.

Neftci, E., Chicca, E., Cook, M., Indiveri, G., & Douglas, R. (2010). State-dependent sensory processing in networks of VLSI spiking neurons. In *International Symposium on Circuits and Systems, ISCAS 2010* (pp. 2789–2792). Piscataway, NJ: IEEE. http://ncs.ethz.ch/pubs/pdf/Neftci_etal10.pdf.

Neftci, E., Chicca, E., Indiveri, G., & Douglas, R. (2011). A systematic method for configuring VLSI networks of spiking neurons. *Neural Computation*, *23*(10), 2457–2497. http://dx.doi.org/10.1162/NECO_a_00182.

Nijmeijer, H. (2001). A dynamical control view on synchronization. *Physica D: Nonlinear Phenomena*, *154*(3–4), 219–228.

Okatan, M., Wilson, M. A., & Brown, E. N. (2005). Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation*, *17*(9), 1927–1961.

Paninski, L., Pillow, J. W., & Simoncelli, E. P. (2004). Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural Computation*, *16*(12), 2533–2561.

Pavasović, A., Andreou, A., & Westgate, C. (1994). Characterization of subthreshold MOS mismatch in transistors for VLSI systems. *Journal of VLSI Signal Processing*, *8*(1), 75–85.

Pecora, L. M., & Carroll, T. L. (1990). Synchronization in chaotic systems. *Phys. Rev. Lett.*, *64*(8), 821–824.

Renart, A., Brunel, N., & Wang, X. (2003). Mean field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 431–490). Boca Raton, FL: Chapman and Hall.

Rossant, C., Goodman, D.F.M., Platkiewicz, J., & Brette, R. (2010). Automatic fitting of spiking neuron models to electrophysiological recordings. *Frontiers in Neuroinformatics*, 1–14.

Russell, A., Orchard, G., & Etienne-Cummings, R. (2007). Configuring of spiking central pattern generator networks for bipedal walking using genetic algorithms. In *International Symposium on Circuits and Systems, 2007* (pp. 1525–1528). Piscataway, NJ: IEEE.

Schenk, O., & Gärtner, K. (2004). Solving unsymmetric sparse systems of linear equations with PARDISO* 1. *Future Generation Computer Systems*, *20*(3), 475–487.

Sheik, S., Stefanini, F., Neftci, E., Chicca, E., & Indiveri, G. (2011). Systematic configuration and automatic tuning of neuromorphic systems. In *International Symposium on Circuits and Systems, ISCAS 2011* (pp. 873–876). Piscataway, NJ: IEEE. http://ncs.ethz.ch/pubs/pdf/Sheik_etal11.pdf.

Shriki, O., Hansel, D., & Sompolinsky, H. (2003). Rate models for conductance-based cortical neuronal networks. *Neural Computation*, *15*(8), 1809–1841.

Stephan, K., Harrison, L., Kiebel, S., David, O., Penny, W., & Friston, K. (2007). Dynamic causal models of neural system dynamics: Current state and future extensions. *Journal of Biosciences*, *32*(1), 129–144.

Toth, B. (2011). *Computational methods for parameter estimation in nonlinear models*. Unpublished doctoral dissertation, University of California, San Diego.

Toth, B., Kostuk, M., Meliza, C. D., Margoliash, D., & Abarbanel, H.D.I. (2011). Dynamical estimation of neuron and network properties I: Variational methods. *Biological Cybernetics*, *105*, 217–237.

Wächter, A., & Biegler, L. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, *106*(1), 25–57.

Wang, X. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *J. Neurosci.*, *19*(21), 9587–9603.

Wang, X. J., & Buzsáki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci*, *16*, 6402–6413.

Wilson, H., & Cowan, J. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*, 1–23.

Wilson, H., & Cowan, J. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Biological Cybernetics*, *13*(2), 55–80.

Yuille, A. L., & Geiger, D. (2003). Winner-take-all networks. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 1228–1231). Cambridge, MA: MIT Press.

Yuille, A. L., & Grzywacz, N. M. (1989). A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Computation*, *1*(3), 334–347.

Received March 30, 2011; accepted December 20, 2011.