

Self-sustained activity in Attractor Networks using Neuromorphic VLSI

Patrick Camilleri, *Student member, IEEE*, Massimiliano Giulioni,
Maurizio Mattia, Jochen Braun, and Paolo Del Giudice

Abstract—We describe and demonstrate the implementation of attractor neural network dynamics in analog VLSI chips [1]. The on-chip network is composed of an excitatory and an inhibitory population of recurrently connected linear integrate-and-fire neurons. Besides from the recurrent input these two populations receive external input in the form of spike trains from an Address-Event-Representation (AER) based system. External AER input stimulates the attractor network and provides also an adequate background activity for the on-chip populations. We use the mean-field approximation of a model attractor neural network to identify regions of parameter space allowing for attractor states, matching hardware constraints. Consistency between theoretical predictions and the observed collective behaviour of the network on chip is checked using the ‘effective transfer function’ (ETF) [2]. We demonstrate that the silicon network can support two equilibrium states of sustained firing activity that are attractors of the dynamics, and that external stimulation can provoke a transition from the lower to the higher state.

I. INTRODUCTION

Neuromorphic chips, purporting to emulate the principles of information processing in the nervous system, have been largely devoted to duplicate in silicon the operation of sensory systems (such as retina [3] or cochlea [4]), and sometimes to implement simple, general purpose computational elements supposedly at work in a variety of neural circuits (such as winner-take-all networks – [5][6]). In many instances, the chosen network architecture is either essentially feedforward [7], or it includes simple feedback mechanisms, as in winner-take-all or Central Pattern Generator (CPG) networks [8]. In the present work we take a step towards silicon implementation of recurrent neural networks with massive feedback, exhibiting attractor behavior. Our main motivation is the belief that attractor networks should be considered as key building blocks of systems including, downstream a possibly neuromorphic sensory system involving complex processing stages, for example effecting a classification of the sensory input or accumulating information about it for a decision to be taken [9]. It has long been recognized that for recurrent networks with high levels of feedback the strength of synaptic connections can be chosen such that the network can store and retrieve prescribed patterns of collective activation as ‘memories’ [10] [11]. Given the initial state of the network, implementing an external stimulus, the network

dynamics relax to the closest fixed point attractor (stored pattern), up to small fluctuations: the network works as an ‘associative memory’, retrieving a prototypical memorized representation for a whole class of stimuli which define the ‘basin of attraction’. If a stimulus is applied and then released, the attractor property of the stored patterns allows the network to sustain a persistent activity pattern which is selective for the stimulus (if it is close enough to a stored memory) and stable in its absence. The network behaves essentially as a bistable system, with two stationary states of low and elevated firing activity, to be associated with the ‘spontaneous’ activity state and a selective state triggered by the stimulus. The above properties make attractor networks of spiking neurons especially suited to provide a dynamic correlate of the persistent neural activity observed in cortex (for example, but not only, in infero-temporal cortex [12] and in prefrontal cortex [13]) in tasks requiring information about a stimulus to be held active in working memory after the stimulus has been removed, for later use in the task. Standard examples include Delayed Match-to-Sample (DMS) tasks [14], in which the subject is required to report if a briefly shown sample image is the same as a match image shown after a delay, or Pair Association tasks [15], in which one of two images shown after the delay has to be chosen, according to a prescribed correspondence to the one shown before the delay. Attractor models have been developed and improved to account for a wide array of experimental evidence related to working memory. Also, it is increasingly becoming clear that the dynamic scheme has a wider scope. Models based on bistable or multistable networks have been proposed as theoretical underpinnings for understanding perceptual decision mechanisms and processes of information integration [16] [11], as well as multi-stable perception and binocular rivalry [17]. It so appears that attractor networks could be considered as general-purpose processing elements, worth the effort to implement them in silicon, in view of complex neuromorphic systems. In the present work we do not consider the unsupervised buildup of stimulus-driven synaptic modifications leading the network to support attractor states, but assign values to the synaptic efficacies such that the resulting neural dynamics exhibit attractor behavior, and check its match with theoretical predictions (though a specific form of Hebbian plasticity is implemented in the chip, and will be used to study the dynamic generation of attractor states in future work).

P. Camilleri and J. Braun are with the Department of Cognitive Biology, Otto-von-Guericke University, Leipziger Str. 44 / Haus 91, 39120 Magdeburg, Germany (email: patrick.camilleri@ovgu.de).

M. Giulioni, M. Mattia, and P. Del Giudice are with the Department of Technologies and Health, Istituto Superiore di Sanità, V.le Regina Elena 299, 00161 Rome, Italy (email: massimiliano.giulioni@iss.infn.it).

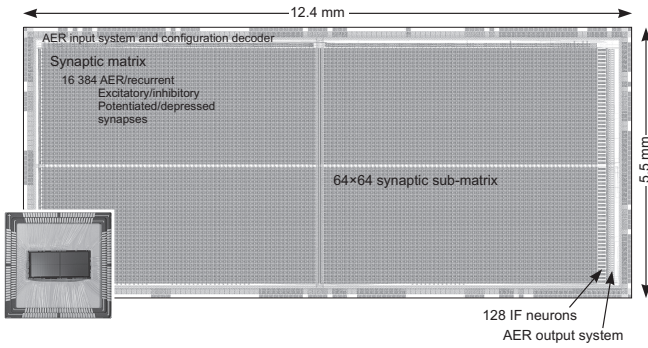


Fig. 1. Chip layout and photo. Chip built using a 0.35 μm AMS CMOS process and has an area of approximately 70 mm^2 .

II. CHIP ARCHITECTURE AND MAIN FEATURES

The chip [1] [18] is an analog VLSI implementation of a network of 128 integrate-and-fire neurons with linear decay (VLSI-IF neurons [19]) and a total of 16,384 bistable plastic synapses, such that each ‘dendritic’ tree of each neuron is made up of 128 synapses. Every synapse can assume one of two possible weight states, a potentiated and a depressed state. Input to each synapse can arrive in the shape of digital spikes from an external Address Event Representation (AER [20] [21]) interface, or directly from any other neuron located on the chip. Arbitrary on-chip synaptic recurrent connectivity can be set, up to an all-to-all. Even though it is possible to achieve a recurrent configuration by means of the off-chip AER infrastructure [7], having the local recurrent synapses gives us the possibility of making the best use out of the available AER bandwidth, at the price of a larger silicon area. Every synapse can be configured to be either excitatory or inhibitory. When configured as excitatory, the synapses also inherit plastic properties, which for the purpose of this experiment were not needed. AER based communication is handled through the PCI-AER board [22] [23]. The *sequencer* of the PCI-AER board allows us to stimulate the hardware system by sending spike trains to the chip generated by the neuron populations simulated in MATLAB. On the input side, the *monitor* of the PCI-AER board acquires AER spikes coming from the chip. To measure the *effective transfer function* (see Section V) we also make use of the *mapper* feature of the PCI-AER board in order to have fast one-to-many AER connections. Exploiting the *sequencer* together with the *mapper* we achieve a bandwidth of up to 0.7 Mega-spikes per second.

As regards to the chip physical details the chip was designed using an AMS CMOS 0.35 μm process, has an area of approximately 70 mm^2 and is housed in a 256-pin PGA package [1].

III. NETWORK ARCHITECTURE

The flexibility of the synaptic matrix allows us to implement different network architectures. In what follows we will refer to the network shown in Figure 2. An on-chip excitatory population (E_{chip}) composed of 50 neurons and an on-chip inhibitory population (I_{chip}) composed of 28

neurons are connected among themselves and both receive external stimuli via the AER bus from three populations ($E1_{\text{pc}}$, $E2_{\text{pc}}$ and I_{pc}) simulated on a standard PC. Intra and inter-populations connectivity levels c are reported in the figure. c stands for the probability that each neuron forms a direct synaptic contact with any other neuron of the target population. For instance, the dendritic tree of each neuron of E_{chip} is made up of $0.25 \cdot 50 \approx 13$ recurrent synapses, $0.21 \cdot 28 \approx 6$ synaptic connections receiving spikes from neurons of I_{chip} , and an additional 70 external AER synapses. Among the AER synapses, 50 accept spikes from the $E1_{\text{pc}}$ excitatory neurons and the remaining 20 receive inputs from the I_{pc} inhibitory neurons. Neuronal populations simulated on the PC are intended to provide both desired stimuli and an adequate background activity for the on-chip populations.

In Figure 2 we also report, for each connection, the synaptic efficacy value J as a fraction of the neuronal dynamic range $\theta - H$, where θ is the neuronal firing threshold and H is the reset potential. Other relevant network parameters are the neuronal leakage current β equal to $35(\theta - H) \text{ s}^{-1}$ and the absolute refractory period $\tau_{\text{arp}} = 2.7 \text{ ms}$ of the neuron, values valid both for neurons belonging to E_{chip} and to I_{chip} . To implement such a network in hardware we first configure the on-chip synaptic matrix and then run a calibration procedure to set the chip bias levels for the VLSI neurons and synapses such that they correspond to the theoretical values of J , β and τ_{arp} . Setting the synaptic matrix presents no difficulties since the operation is based on a simple digital protocol handled by a dedicated microcontroller. Setting biases is a more demanding task which is described in detail in the next section.




Fig. 2. Architecture of the network implemented on chip: E_{chip} consists of 50 excitatory neurons and I_{chip} consists of 28 inhibitory neurons. For each connection, the synaptic efficacy (J) and the connectivity level (c) is specified. Apart from the on-chip connections, each neuron of E_{chip} has 50 AER excitatory synapses addressed from the population $E1_{\text{pc}}$ and 20 AER inhibitory synapses receiving spikes from I_{pc} . The AER part of each dendritic tree of neurons in I_{chip} consists of 50 excitatory synapses accepting spikes from $E2_{\text{pc}}$. $E1_{\text{pc}}$, $E2_{\text{pc}}$ and I_{pc} are simulated on a standard PC and consist of 2500, 1400, and 1000 neurons respectively, each connected to one of the on-chip AER synapses.

IV. SETTING THE CHIP BIAS VOLTAGES

Since the neuron and synapse circuits are based on a sub-threshold circuit implementation, the circuits are sensitive to semiconductor process variations, temperature fluctuations and power supply voltage drops. This results in on-chip mismatch which leads to a certain distribution of the neural and synaptic parameters, such as synaptic weight J , neuron leakage current β , and refractory period τ_{arp} . Moreover, parameters are slightly coupled together creating a challenge in finding the right chip biases to make them match with the actual theoretical values of a particular experiment we want to reproduce in VLSI.

We start off from the equation of the mean emission rate of a VLSI neuron as a function of the mean and variance of the input current defined by [19]:

$$\nu \equiv \Phi(\mu, \sigma) = \left[\tau_{\text{arp}} + \frac{\sigma^2}{2\mu^2} \left(e^{-\frac{2\mu(\theta-H)}{\sigma^2}} - 1 + \frac{2\mu(\theta-H)}{\sigma^2} \right) \right]^{-1}$$

If one assumes that the mean number of afferent connections is large, the mean synaptic efficacy relative to the neuronal dynamic range $\theta-H$ is small, and that the emission times of the various neurons are uncorrelated [19], then for the excitatory populations:

$$\begin{aligned} \mu &= c_{ee}N_eJ_{ee}\nu_e + c_{ei}N_iJ_{ei}\nu_i + \mu_{\text{ext}} - \beta \\ \sigma^2 &= c_{ee}N_eJ_{ee}^2\nu_e + c_{ei}N_iJ_{ei}^2\nu_i + \sigma_{\text{ext}}^2 \end{aligned}$$

where the terms depending on ν_e and ν_i are due to the recurrent connections, and the offset terms $\mu_{\text{ext}} - \beta$ and σ_{ext}^2 are due to the external spikes and the constant leakage term β of the neuron. In the above equations c represents the connectivity as a fraction of the population, N the population size, and J the synaptic efficacy as a fraction of the neuronal dynamic range $\theta - H$.

Instead of checking the value of each chip parameter at any point of the parameter space, we focus on specific network conditions suggested by the mean-field theory, and the strategy we chose to correctly tune the biases is to match the average of multiple single-neuron transfer functions [19] with their theoretical counterpart. To measure the neuron transfer functions, synapses of a neuron are configured to receive spikes from an external source and the neuron of interest is characterized by stimulating it using gaussian distributed spike trains at a given frequency ν_{in} . The output firing rate Φ of every neuron in the population of interest is measured by means of the PCI-AER interface and the average computed.

In order to determine the chip analog biases necessary to produce the desired theoretical parameters, we first start by getting a rough idea of the actual chip parameter values with the help of an oscilloscope and a few MATLAB scripts analogous to those already described in [24]. We then proceed by considering the network depicted in Figure 2 as if it consisted of two separate populations with no recurrent and inter-population connections. Once the populations are considered as being disconnected from their environment we

can start measuring the individual neuron transfer functions. We commence by setting all the chip synapses to be AER, excitatory, and depressed which in turn enables us to find the excitatory depressed efficacy J_{exc} , the constant leakage term β of the neuron and the absolute refractory period τ_{arp} . The on-chip EPSPs and IPSPs are due to a constant current having an amplitude defined by J_{exc} which is in turn applied to the target neuron for approximately 2.5 ms. Note that when the AER synapses are reconfigured as recurrent, the value of J_{exc} just found is retained. J_{exc} found here is the on-chip efficacy value used for the recurrent connections of the excitatory population, the connection strength between the excitatory and the inhibitory populations, and the strength between the external excitatory populations and the chip populations (see Figure 2). The last step involves configuring a subset of the synapses as inhibitory which enables us to tweak the J_{inh} bias to set the inhibitory connection strengths used for the recurrent inhibitory connections and for the efficacy of the inhibitory to excitatory connections. The result of this last step is depicted in Figure 3. Iterating over these two steps enables us to obtain an accurate correspondence between the theoretical and measured transfer functions [19]. Once the main parameters (J_{exc} , J_{inh} , β , and τ_{arp}) are determined, we make sure that they really are the correct ones by plugging them in the actual network represented in Figure 2 and measuring the *effective transfer function* of the network (see Section V). The reason we go through all these steps to find the chip biases is to be able to reduce the number of simultaneous bias settings that we would otherwise have to find with the help of the *effective transfer function* alone.




Fig. 3. Neuron transfer function $\Phi(\mu, \sigma)$ with a subset of the synapses configured as inhibitory. The rest of the synapses are set to be excitatory and depressed.

V. EFFECTIVE TRANSFER FUNCTION

For a given set of parameters, the equilibrium states (fixed points) of the network dynamics are found as the studies of the self-consistency equation $\nu = \Phi(\nu(\mu, \sigma^2))$ (supplemented by linear stability analysis). For p interacting

populations, the equilibrium states are formed by solving the system $\underline{\nu} = \Phi(\underline{\nu})$. Since it can be difficult to get an intuitive picture of such a multi-dimensional problem, an approximate approach has been proposed in [2], which allows to focus on a subset of the population involved and compute an *effective transfer function*. The essence of the method is an iterative procedure by which at each step one fixes the ν of the populations in focus and uses them as parameters in the self-consistency equations for the remaining population. For example, focusing on population no. 1 and fixing $\nu_1 = \bar{\nu}_1$, the rest of the network adapts to ν_1 reaching a global equilibrium state $\nu_2^*(\bar{\nu}_1), \dots, \nu_p^*(\bar{\nu}_1)$:

$$\begin{aligned} \nu_2^* &= \Phi_2(\bar{\nu}_1, \nu_2^*, \dots, \nu_p^*) \\ &\vdots \\ \nu_p^* &= \Phi_p(\bar{\nu}_1, \nu_2^*, \dots, \nu_p^*) \end{aligned}$$

The new state $(\nu_2^*, \dots, \nu_p^*)$ drives population no. 1 to a new rate

$$\bar{\nu}'_1 = \Phi_1(\bar{\nu}_1, \nu_2^*, \dots, \nu_p^*) \equiv \Phi_{\text{eff}}(\bar{\nu}_1)$$

effectively reducing the mean-field formulation to a 1-dimensional problem for the considered population, which embodies the full effect of the feedback among the other populations. Therefore, after having identified appropriate parameters through the analysis of the single-neuron transfer function, a sensible strategy is to rely on the *effective transfer function* to adjust synaptic efficacies suited to support attractor states for the whole excitatory-inhibitory on-chip network.

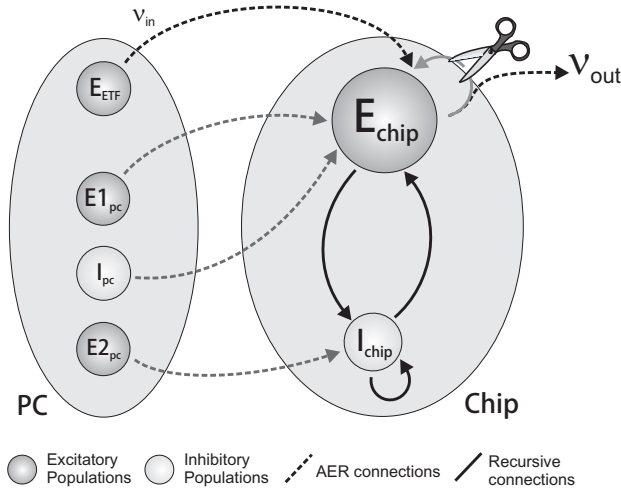


Fig. 4. Effective transfer function architecture.

To measure the *effective transfer function*, we consider a modified version of our network (see Figure 4): modifications consists in the cutting of the recursive connections of E_{chip} and introducing a new external AER excitatory population E_{ETF} . Each neuron of E_{chip} has its alter-ego in E_{ETF} and the connectivity between these two populations exactly reproduces the severed recurrent connections. If in our original network, neuron A of E_{chip} is pre-synaptic to neurons B and

C of E_{chip} , an alter-ego of neuron A will now be simulated in the external population E_{ETF} and will be connected via AER to neurons B and C of E_{chip} . On-chip, starting from the synaptic configuration implementing the architecture shown in Figure 2, we simply turn recursive synapses of E_{chip} into AER synapses addressed from the corresponding neurons of E_{ETF} . All the other synaptic connections are left untouched. The necessary AER bandwidth is assured by the use of the *mapper* of the PCI-AER board designed to provide fast one-to-many connections. A spike emitted by E_{ETF} is physically generated by the *sequencer* and passed onto the *mapper* which in turn generates at a hardware level, a burst of spikes sent to the target neurons.

Population E_{ETF} is simulated on a standard PC such that one can decide its firing rate: the input frequency ν_{in} to our black box system; the output of the system is the mean firing rate ν_{out} of neurons in E_{chip} . Performing a sweep of ν_{in} one obtains the *effective transfer function* of the system (see Figure 5). This curve represents, to a certain degree of approximation [2], a ‘static’ version of the trajectory the system will follow in time if the excitatory recurrent connections are restored. The basic idea is that by disconnecting the recurrent connections, one prevents the system from autonomously evolving in time: stimulating the system with a constant ν_{in} , the network adapts and produces a certain ν_{out} which can be fed back in as a new ν_{in} in the next stimulation. Iterating this procedure one explores, step by step, various transient states of the network dynamics. The graph in Figure 5 shows that the stable points of the network dynamics are at about 0.5 Hz and 160 Hz. The second intersection of the *effective transfer function* with the line $\nu_{\text{in}} = \nu_{\text{out}}$ at approximately 40 Hz is an unstable point of the system dynamics (see [19]) and represents the barrier the network has to cross to jump from one stable state to the other. Another fact we would like to stress is that Figure 5 reports the measure of the population activity and shows that, even if parameters mismatch affects chip activity at a neuronal level, at a population level the hardware behaves in agreement with the theoretical mean-field prediction. This is reasonable if one considers the amount of interaction among neurons and the averaging effect taking place on the dendritic trees composed of up to about 90 synapses.

To obtain the *effective transfer function* for a given ν_{in} we stimulate the system for ten seconds while monitoring ν_{out} with the PCI-AER board. Neurons belonging to E_{ETF} emit spikes according to a gaussian ISI (Inter Spike Interval) distribution, centered on ν_{in} with a standard deviation equal to 10% of ν_{in} . During all the stimulations, populations $E1_{\text{pc}}$, $E2_{\text{pc}}$ and I_{pc} maintain a constant firing rate of 2 Hz, 3.9 Hz, and 7 Hz respectively, the same rates they have during the experiment explained in the next section.

VI. ATTRACTOR

Summing up: starting from a good point suggested by the mean-field theory, we chose a set of parameters and found the corresponding biases on the chip. Measuring the effective transfer function we completed the iterative process for

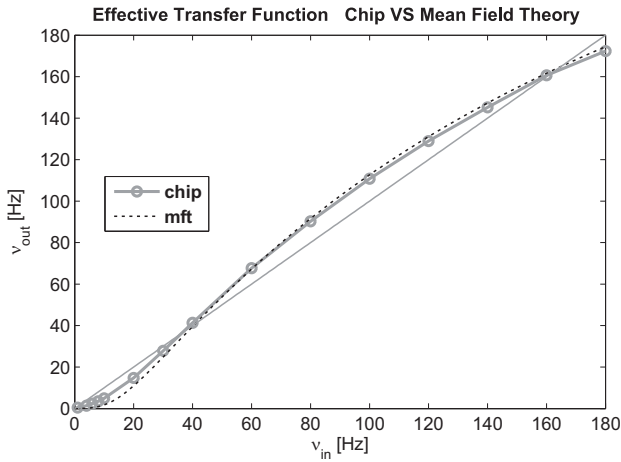


Fig. 5. Effective transfer function. On the x -axis the input frequency, i.e. the mean firing rate of E_{ETF} , on the y -axis the system output, i.e. the mean firing rate of E_{chip} . The grey solid line reports data measured from the chip, the black dashed line results from the mean-field theory. The black solid line is the line $\nu_{out} = \nu_{in}$. Intersections between the diagonal and the effective transfer function indicate fixed points of the network dynamics.

setting the biases and checked the behavior of the chip at a population level. We now restore the recursive connections of E_{chip} and run a stimulation protocol to demonstrate that the network has two different states of activity. The stimulation protocol is divided into three phases during which everything remains unchanged except for the mean frequency (ν_{E1}) of the spikes emitted by $E1_{pc}$. In the first phase, lasting 1 second, the level of external stimulation is low ($\nu_{E1} = 2$ Hz). In the second phase lasting one second ν_{E1} is increased by a factor of 2.4. During the third phase the mean frequency of $E1_{pc}$ is reduced again to its original value. The mean firing rates of $E2_{pc}$ and I_{pc} are constant at 3.9 and 7 Hz throughout the 3 phases. Figure 6 reports the frequency profile of the populations E_{chip} and I_{chip} (solid line) during the stimulation: in black the excitatory population, in grey the inhibitory one. The increase in external frequency provided the network the necessary energy to jump from the lower stable state, where the main contribution to the network activity is given by the external AER populations, to the upper stable state where the mean firing rate of E_{chip} is about 160 Hz in agreement with the mean-field theory and with the effective transfer function prediction. Supported by local reverberations the network remains in this upper state during the entire duration of the third phase showing a persistent stable activity triggered by the increase in the external stimulus. In Figure 6 the theoretical behavior predicted by the mean field theory is reported in dashed lines. A fundamental property of an attractor is its ability to recruit all neurons in the “cell assembly” even when the input stimulus is corrupted [11]. To test this property we reduce the number of stimulated neurons N_{stim} of E_{chip} receiving a stronger stimulus from the external population $E1_{pc}$ during the second phase of the protocol described above. In Figure 7, $N_{stim} = 26$ for the upper graph and $N_{stim} = 20$ for the lower graph. Black lines show the mean frequency profile of the N_{stim} neurons, while in grey

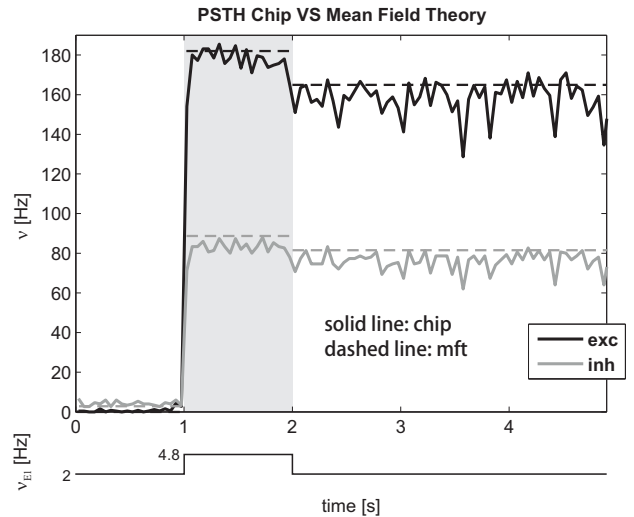


Fig. 6. Profile frequency of the two on chip populations E_{chip} (black) and I_{chip} (grey). From 1 s to 2 s the network receives an increased external stimulus. Solid lines show the frequency profile measured from the chip, dashed lines the corresponding mean-field prediction. Below the graph the frequency profile of the external population $E1_{pc}$ is reported.

the mean frequency profile of the non-stimulated neurons of E_{chip} is reported. When the stimulated (black) neurons are able to bring the non-stimulated (grey) neurons above the barrier defined by the unstable fixed point of the dynamics, the entire network undergoes the transition to the upper stable state and the frequency fluctuations of the two groups of neurons become more strongly correlated, demonstrating the attractor nature of the dynamics. Results reported in Figure 7 are only two examples: since the nature of the dynamics is stochastic and depends on the particular instances of the spike trains stimulating the network, accumulating statistics, one can show that reducing N_{stim} the probability of jumping to the upper stable states decreases and that this decrease depends on the time duration of the second phase of the protocol. A more detailed characterization of the network properties will be published in a future work.

VII. CONCLUSIONS

We demonstrate the operation of a VLSI recurrent neural network of spiking neurons supporting discrete metastable attractor states of very low activity and elevated activity, with transitions between the two allowed for sufficiently strong stimulation. The silicon network behavior matches well the one predicted by the mean-field theory in terms of an effective transfer function. The attractor property of such states is confirmed by stimulating a subset of neurons and showing that the collective dynamics are quickly recruiting the non stimulated neurons thanks to extensive feedback.

This is, to our knowledge, the first demonstration of a silicon, recurrent network exhibiting discrete attractor states (but see [25] for the VLSI implementation of a continuous attractor model). To the extent that attractor models can be considered as biologically relevant components of a variety of computational scenarios, as argued in the introduction, this

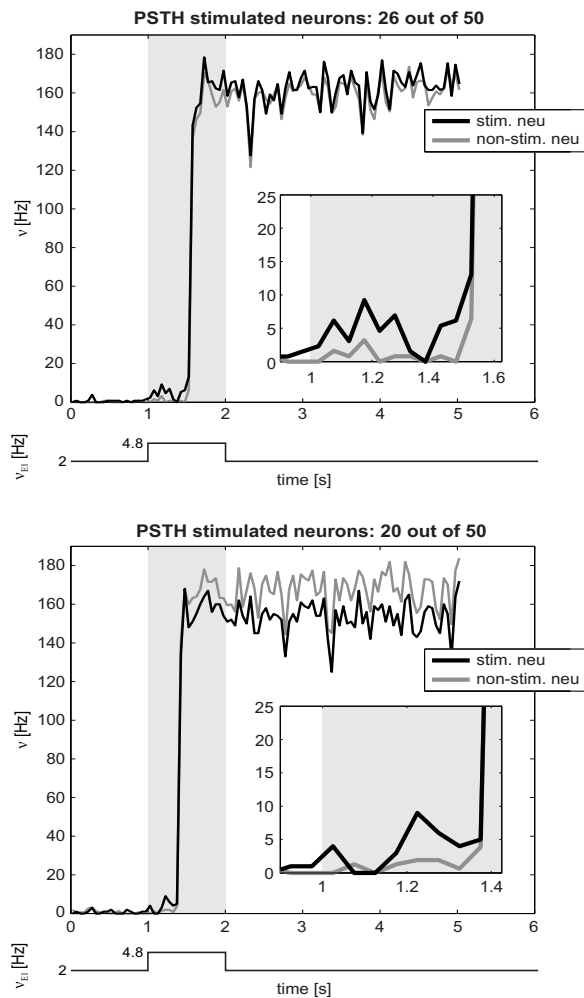


Fig. 7. Testing the ability of recruiting neurons. All the curves refer to neurons of E_{chip} . In black the mean frequency profile of the N_{stim} neurons receiving an increased stimulus from $E1_{PC}$ during the second phase of the protocol; in grey the mean frequency profile of the non-stimulated neurons in E_{chip} . In the left panel $N_{stim} = 26$, in the right panel $N_{stim} = 20$. In the right panel the small difference between the frequency asymptotic levels is due to VLSI mismatch. Below each graph the frequency profile of the external population $E1_{PC}$ is reported.

constitutes a significant step towards biologically-inspired information processing systems.

On the technical side, it is also rewarding to verify that the approach taken in the chip design, together with the PCI-AER programmable interface, ensure easy and flexible configuration of the synaptic connectivity, high-level and user-friendly parameters setting and chip interaction with a synthetic environment emulating additional simulated neural populations.

REFERENCES

[1] M. Giulioni, P. Camilleri, V. Dante, D. Badoni, G. Indiveri, J. Braun, and P. Del Giudice. A VLSI network of spiking neurons with plastic fully configurable stop-learning synapses. In *Proc. IEEE International Conference on Electronics Circuits and Systems*, pages 678–681, 2008.

[2] M. Mascaro and D. J. Amit. Effective neural response function for collective population states. *Network: Computation in Neural Systems*, 10:251–373, 1999.

[3] P. Lichtsteiner, C. Posch, and T. Delbruck. A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid State Circuits*, 43(2):566–576, 2008.

[4] V. Chan, S. C. Liu, and A. van Schaik. AER EAR: A matched silicon cochlea pair with address event representation interface. *IEEE Transactions on Circuits and Systems I: Special Issue on Smart Sensors*, 54(1):48–59, 2007.

[5] G. Indiveri. A current-mode analog hysteretic winner-take-all network, with excitatory and inhibitory coupling. *Analog Integrated Circuits and Signal Processing*, 28(3):279–291, 2001.

[6] J. P. Abrahamsen, P. Haffiger, and T. S. Lande. A time domain winner-take-all network of integrate-and-fire neurons. In *Proc. IEEE International Symposium on Circuits and Systems ISCAS04*, pages 361–364, 2004.

[7] S. Mitra, S. Fusi, and G. Indiveri. Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. *IEEE Transactions on Biomedical Circuits and Systems*, 3(1):32–42, 2009.

[8] R. J. Vogelstein, F. Tenore, L. Guevremont, R. Etienne-Cummings, and V. K. Mushahwar. A silicon central pattern generator controls locomotion in vivo. *IEEE Transactions on Biomedical Circuits and Systems*, 2(3):212–222, 2008.

[9] D. Marti, G. Deco, M. Mattia, G. Gigante, and P. Del Giudice. A fluctuation-driven mechanism for slow decision processes in reverberant networks. *PLoS One*, 3:e2534, 2008.

[10] D. J. Amit. The hebbian paradigm reintegrated, local reverberations as internal representation. *Behavioral and Brain Science*, 18:617–657, 1995.

[11] D. J. Amit. *Modeling Brain Function*. Cambridge University Press, 1989.

[12] J. M. Fuster and J. P. Jervey. Inferotemporal neurons distinguish and retain behaviorally relevant features of visula stimuli. *Science*, 212(4497):952–955, 1981.

[13] J. M. Fuster and G. E. Alexander. Neuron activity related to short-term memory. *Science*, 173(3997):652–654, 1971.

[14] Y. Miyashita. Neuronal correlates of visual associative long term-memory in the primate temporal cortex. *Nature*, 335(6193):817–820, 1988.

[15] K. Sakai and Y. Miyashita. Neural organization of the long-term memory of paired associated. *Nature*, 354(6349):152–155, 1991.

[16] X. J. Wang. Decision making in recurrent neuronal circuits. *Neuron*, 60(2):215–234, 2008.

[17] G. Gigante, M. Mattia, J. Braun, and P. Del Giudice. Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Comp. Biol.* Accepted, 2009.

[18] M. Giulioni. *Networks of spiking neurons and plastic synapses: implementation and control*. PhD thesis, University of Rome “Tor Vergata”, 2008.

[19] S. Fusi and M. Mattia. *Neural Computation*, 11:633, 1999.

[20] M. Mahowald. *VLSI analogs of neuronal visual processing: a synthesis of form and function*. PhD thesis, California Institute of Technology, Pasadena, CA, 1992.

[21] K. A. Boahen. Point-to-point connectivity between neuromorphic chips using address-events. *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, 47(5):416–434, 2000.

[22] P. Del Giudice V. Dante and A. M. Whatley. The neuromorphic engineer newsletter. <http://ine-web.org/fileadmin/templates/docs/nme3.pdf>. 2005.

[23] E. Chicca, V. Dante, A.M. Whatley, P. Lichtsteiner, T. Delbruck, G. Indiveri, P. Del Giudice, and R. J. Douglas. Multi-chip pulse based neuromorphic systems: a general communication infrastructure and a specific application example. *IEEE Transactions on Circuits and Systems I*, 54(5):981–993, 2007.

[24] M. Giulioni, M. Pannunzi, D. Badoni, V. Dante, and P. Del Giudice. Classification of overlapping patterns with a configurable analog VLSI neural network of spiking neurons and self-regulating plastic synapses. *Neural Computation*, 21(11):3106–3129, 2009.

[25] T. M. Massoud and T. K. Horiuchi. A neuromorphic head direction cell system. In *Proc. IEEE International Symposium on Circuits and Systems ISCAS09*, 2009.