

State-dependent sensory processing in networks of VLSI spiking neurons

Emre Neftci, Elisabetta Chicca, Matthew Cook, Giacomo Indiveri, Rodney Douglas
 Institute of Neuroinformatics
 University of Zurich and ETH Zurich
 Email: emre@ini.phys.ethz.ch

Abstract—An increasing number of research groups develop dedicated hybrid analog/digital very large scale integration (VLSI) devices implementing hundreds of spiking neurons with bio-physically realistic dynamics.

However, despite the significant progress in their design, there is still little insight in translating circuitry of neural assemblies into desired (non-trivial) function.

In this work, we propose to use neural circuits implementing the soft Winner-Take-All (WTA) function. By showing that recurrently connected instances of them can have persistent activity states, which can be used as a form of working memory, we argue that such circuits can perform state-dependent computation. We demonstrate such a network in a distributed neuromorphic system consisting of two multi-neuron chips implementing soft WTA, stimulated by an event-based vision sensor. The resulting network is able to track and remember the position of a localized stimulus along a trajectory previously encoded in the system.

I. INTRODUCTION

The brain is able to process its sensory modalities and achieve complex behavior with great performance and robustness, requiring very little power. This has motivated several scientists to focus their efforts in the design of low-power hybrid silicon analog/digital circuits implementing neurons and synapses [1], which operate in parallel and respond in real-time. If such devices could be deployed and interconnected in a way to perform the function desired by the user, they could become ideal computational modules for event-driven sensory systems including silicon cochleae and/or retinæ [2], [3].

There is no doubt that the remarkable capabilities of the brain emerge from the numerous interactions between its different subunits. In fact, it has been argued that a good candidate model for a canonical micro-circuit, potentially used as a general purpose cortical computational unit is the soft Winner-Take-All (WTA) circuit [4], and that recurrently connected instances of them is able to perform any desired computation [5].

Complex behavior requiring the sequential completion of a number of subtasks [6] can be performed by making use of persistent activity states, which act as a form of working memory. In Sec. II, we demonstrate the existence of persistent activity states in the soft WTA, and argue that they can be used for state-dependent computation. In Sec. III, we present a neuromorphic setup consisting of two recurrently connected multi-neuron chips implementing a soft WTA and using persistent activity states to process the sensory output of an event-based vision sensor in a state-dependent fashion.

II. RECURRENTLY COUPLED LINEAR THRESHOLD UNITS FOR STUDYING SPIKING NEURON ACTIVITY

The theoretical analysis of networks of spiking neurons can be carried out using models based on Linear Threshold Units (LTUs), which represent the *mean firing rate* of biological neurons [7]–[9]. In a soft WTA of the type described in Fig. 1(a), only a localized population of neurons is active at any given time (often described as a “bump” of activity [?]), such that the active population behaves approximatively as a single LTU [7]. Thus, for the analysis in this section, we can model an active population and its recurrent couplings in the soft WTA and the global inhibition with a pair of LTUs. In this model, the excitatory units excite themselves with weight w_e and their respective inhibitory unit with weight w_{ie} . When the inhibitory units are active, they inhibit their respective excitatory units with weight w_{ie} (see Fig. 1(a)).

Our goal is to associate two features of a common stimulus, each of which will be represented by one soft WTA (see Sec. III). Therefore we will study the persistent activity states in two pairs of LTUs whose excitatory units are bi-directionally coupled with weight γ . The dynamics of both pairs of LTUs can then be expressed as:

$$\begin{aligned}\tau \dot{x}_{e1} &= -x_{e1} + \sigma(w_e x_{e1} - w_{ie} x_{i1} + \xi \gamma x_{e2} - T + I_1) \\ \tau \dot{x}_{i1} &= -x_{i1} + \sigma(w_{ie} x_{e1} - T) \\ \tau \dot{x}_{e2} &= -x_{e2} + \sigma(w_e x_{e2} - w_{ie} x_{i2} + \xi \gamma x_{e1} - T + I_2) \\ \tau \dot{x}_{i2} &= -x_{i2} + \sigma(w_{ie} x_{i2} - T)\end{aligned}$$

where $\sigma = \max(0, x)$ is a half-wave rectification non-linearity, $\xi \in \{0, 1\}$ is a binary variable specifying the state of the coupling between both pairs of LTUs, T is the threshold of the LTU, I_1 and I_2 are the time-dependant external inputs to the respective excitatory units and τ is the time constant of the LTUs.

A. Existence of a persistent activity state

We study the range of parameters in which the persistent activity state exists, even in the absence of input ($I_1 = 0, I_2 = 0$). It can be easily shown that the equilibrium state and the largest eigenvalue of the Jacobian of the single-pair system are identical to those of the double-pair system. Therefore the stability analysis may be reduced to the case $x_{e1} = x_{e2} \equiv x_e$, $x_{i1} = x_{i2} \equiv x_i$.

The *nullclines*, which are the curves defined by $\dot{x}_e = 0$ and $\dot{x}_i = 0$, are useful for identifying equilibrium points. In the




Fig. 1. *Analysis of LTU model of recurrently coupled soft WTA* (a) (bottom) shows the architecture of the soft WTA network used in our chips. The excitatory neurons (white) excite their nearest neighbors and the inhibitory neurons (gray) inhibit the excitatory neurons back, leading to a single region of activity in the excitatory neurons. (a) (top) The LTU model of two recurrently coupled soft WTA. (b) Nulllines of the LTU model. The curves indicate when \dot{x}_e or \dot{x}_i change sign and their intersection defines an equilibrium point of the system. From this figure, we see that the intersection occurs only when the slope of \dot{x}_i is larger than the one of \dot{x}_e .

linear region of σ , *i.e.* $\sigma(x) = x$, they can be respectively expressed as:

$$\begin{aligned} \dot{x}_e = 0 &\Leftrightarrow \{(x_i, x_e) \mid x_i = \frac{(w_e + \gamma - 1)x_e - T}{w_{ie}}\} \\ \dot{x}_i = 0 &\Leftrightarrow \{(x_i, x_e) \mid x_i = w_{ei}x_e - T\} \end{aligned}$$

and are depicted in Fig. 1(b). By definition of the nullclines, the intersection of the curves define the equilibrium point, which should be in the $x_e > 0$, $x_i > 0$ quadrant if we want both units to be active. This immediately imposes $w_{ie} > 1$ and $w_e + \xi\gamma - 1 > w_{ei} > 0$. The LTU model is exponentially stable if the excitatory units do not hit their threshold and if the largest eigenvalue of its Jacobian is negative [10]. A straightforward calculation shows that this is true when $1 - (w_e + \xi\gamma) + w_{ie}w_{ei} > 0$, meaning that the system gain must be finite. An intuitive explanation to the persistent activity state is the following: in the absence of inhibition, the system would be unstable and its activity would increase exponentially. However, the inhibitory unit, which activates only when it reaches its threshold T , has a gain large enough to keep the excitatory unit under control. This positive and negative feedback act as soft boundaries for the activity of the excitatory LTUs thus creating the attractor state.

We are interested in cases where the persistent activity states exists only when the two pairs are coupled ($\xi = 1$). The condition for no persistent activity state when $\xi = 0$ is therefore $w_e < 1$. On the other hand, to ensure that the persistent activity exists when the two units are connected ($\xi = 1$), γ must verify: $\gamma > 1 - w_e$.

The model presented above has only one active state. It can be extended to incorporate more states by adding an additional excitatory LTU per desired state (see [5] for a detailed presentation). When the values of I_1 and I_2 are chosen such that they are unable to activate a new state while another state is active, unless the given transition is allowed (for example when the excited state is positively biased by




Fig. 2. *Visual stimuli and DVS output*. (a) The target was a dark circle over a white background, initially moving horizontally to the right (gray arrow), then pausing 2s before moving either to the lower-right (green arrow) or to the upper-right (red arrow). (b) Example of the DVS output. The axes respectively represent the X-Y coordinates of the events and the color encodes for time. The scattered events around the main stimulus is due to spontaneous activity in the DVS. (c) and (d): raster plots of the DVS column-wise activity and the row-wise DVS activity sent to the respective multi-neuron chips (see Fig. 3). Because the DVS is designed to respond to temporal contrast, no activity is generated when the target remains stationary ($0s > t > 1s$, $3s > t > 5s$, $t > 7s$).

the active state), then the network is able to perform state-dependent computation [5]. As a result, when the network is in an active state and it is stimulated in a region where $\xi = 0$, its activity will remain "stuck" in its last valid state.

The principles described in this section are used as guidelines for the neuromorphic implementation of a state-dependent processing system, which is described in the next section.

III. APPLICATION TO A MULTI-CHIP SYSTEM WITH VISUAL SENSORY INPUT

A. Description of the neuromorphic setup

The VLSI multi-neuron chips used in this work consist of low-power I&F neurons with dynamic synapses [11], [12]. The chip has been fabricated using a standard AMS 0.35 μm CMOS process, and covers an area of about 10 mm^2 . It contains 124 excitatory neurons with local hard-wired self, 1st, 2nd, 3rd nearest-neighbor recurrent excitatory connections and 4 inhibitory neurons (all-to-all bi-directionally connected to the excitatory neurons). When properly configured, the hard-wired connections implement a soft WTA. Each neuron receives input currents from a row of 32 afferent plastic synapses that use the Address-Event-Representation (AER) [13] to receive spikes.

We can interface the chip to a workstation using dedicated boards, allowing us to stimulate the synapses on the chips, monitor the activity of the neurons [14], and map events from one neuron to a synapse belonging to a neuron on

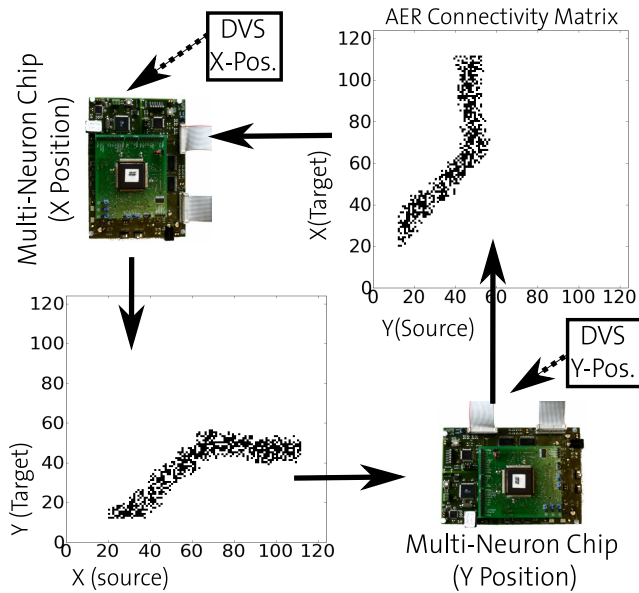


Fig. 3. *AER connectivity between the multi-neuron chips.* One chip is stimulated by the columns of the DVS (Raster plot of Fig. 2(c)), and the other chip is stimulated by the rows of the DVS (Raster plot of Fig. 2(d)). The bi-directional connectivity is represented by the matrices between the two chips and was set during an initialization procedure (see Sec. III). Each black dot in the matrix shows that the particular connection is active. As discussed in Sec. II-A, persistent activity states are created in regions where the two chips are connected. As a result, the location of the activity between the two chips will be constrained by the connections between them.

the same chip and/or on a different chip. Therefore arbitrary connectivity patterns can be implemented in addition to the local hard-wired couplings.

We have implemented the model described in Sec. II with two multi-neuron chips, stimulated by an AER asynchronous 64×64 Dynamic Vision Sensor (DVS) that responds to temporal contrast [15]. The conditions described in Sec. II-A were manually verified. In particular, the chip parameters related to the weights w_e and γ were manually tuned until no persistent activity state was observable when $\xi = 0$.

The DVS was stimulated by an LCD screen presenting a dark circle over a white background (see Fig. 2(a), generated by VisionEgg software [16]).

The column-wise activity of the DVS was retinotopically mapped to the first multi-neuron chip and the row-wise activity was retinotopically mapped to the second multi-neuron chip (see Fig. 2). The connectivity between the two multi-neuron chips was set manually during an initialization procedure where the DVS was stimulated in the absence of recurrent connectivity between the two chips. The response of the two chips was respectively recorded in two matrices A_1 and A_2 representing their $5ms$ -binned firing rates. The neurons of the two chips were then bi-directionally connected together with probability 0.7 if they had fired during the same $5ms$ -time bin. *i.e.* the i, j elements of the connectivity matrix were computed with $M_{ij} = m$ if $[A_1 \cdot A_2]_{ij} > 0$, where $m = 1$ with probability 0.7 and 0 otherwise (see Fig. 3). At the end of this procedure, the system had encoded the relation representing the measured X–Y trajectories of the stimulus.

During the testing phase the network was stimulated with the stimulus used for the initialization procedure, and its response was recorded (see Fig. 4, Top). Because the chips were stimulated in regions where their populations were connected ($\xi = 1$), the activity of the system followed the input, reached its final position and remained in a persistent activity state. However, when the network was stimulated with the incorrect stimulus, as shown in Fig. 4 (Bottom), its activity remained at its last valid location due to the lack of recurrent connectivity between both populations ($\xi = 0$). Depending on the noise and the mismatch of the system, we also observed cases where the activity died out as a result of sudden increase in inhibitory feedback.

IV. DISCUSSION AND CONCLUSION

When the parameters of the neurons were (manually) tuned within the range discussed in Sec. II-A, the recurrent AER connectivity between the two chips gave rise to persistent activity, that was stable even in the absence of external stimuli. In the example shown in this work, the system was able to accept or reject the trajectory of the target, without precise timing of its motion or direction of movement. This example illustrates state-dependent processing of an event-based visual input using a network of recurrently connected VLSI circuits implementing soft WTA.

Although the features related by the network were the X position and the Y position of a moving target, one could easily extract low-level features of the visual stimulus such as orientation bars or Gabors using additional layers of neurons, and apply the same principles discussed in this work.

We observe in Fig. 4 that the activity occasionally resists the drift induced by the stimulus. This is due to the mismatch in the transistors which create persistent activity states which are more or less strong. We also observed cases where the activity died out after it was stimulated in a region where the two chips were not connected ($\xi = 0$). This was due to the effect of mismatch, because of its occurrence in the same location within the array of neurons. These effects reduce the performance of the system, in future work we will include a learning step to our initialization procedure which will automatically adapt the connection probability between the two chips such that it would compensate the mismatch in the transistors.

ACKNOWLEDGMENT

This work was supported by the DAISY (FP6-2005-015803) EU grant, the Swiss National Science Foundation (PMPD2-110298/1), and by the EU ICT Grant ICT-231168-SCANDLE "acoustic SCene ANalysis for Detecting Living Entities". We thank S. Sheik, F. Stefanini and D. Sonnleithner for their help in building the setup, D. Fasnacht for the design of the AER mapper board and the AER monitor/sequencer board, T. Delbrück and P. Lichtsteiner for the vision sensor.

REFERENCES

- [1] C. Mead, *Analog VLSI and Neural Systems*. Reading, MA: Addison-Wesley, 1989.
- [2] P. Lichtsteiner, C. Posch, and T. Delbrück, "A 128×128 120dB 30mW asynchronous vision sensor that responds to relative intensity change," in *2006 IEEE ISSCC Digest of Technical Papers*. IEEE, Feb 2006, pp. 508–509.

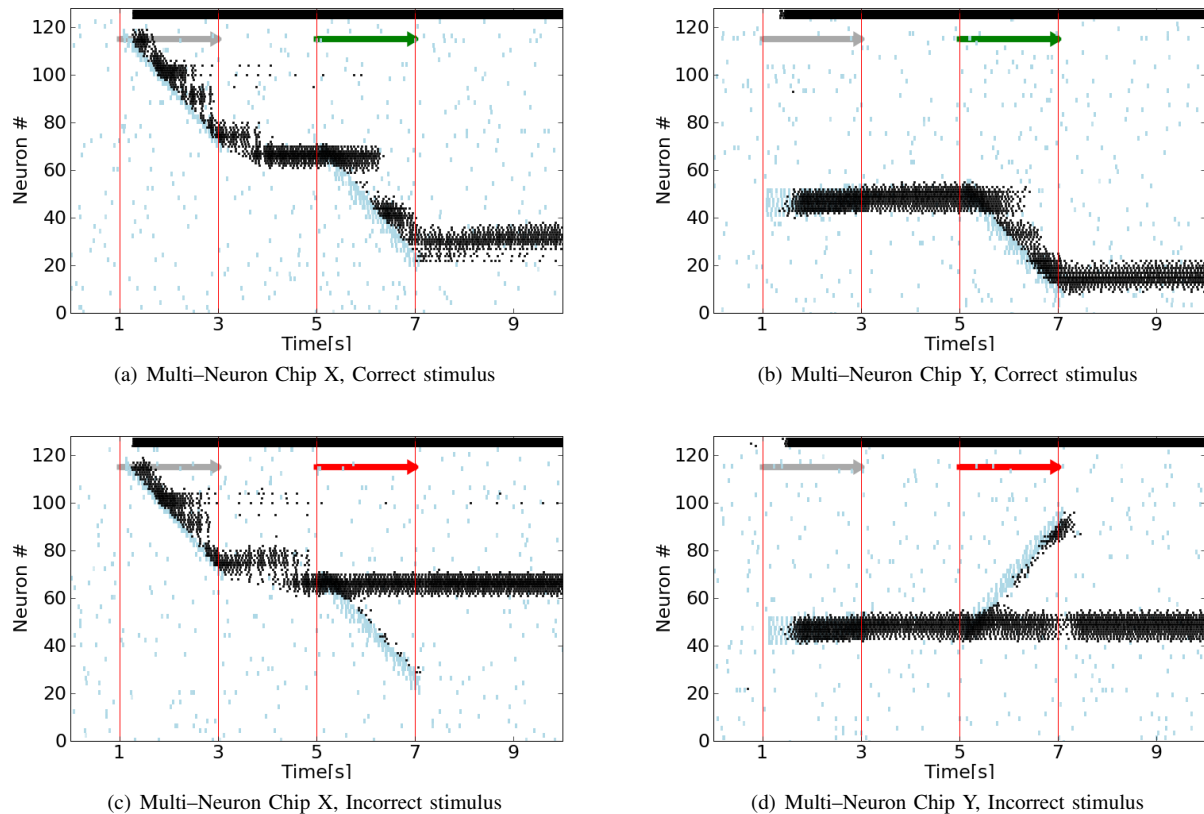


Fig. 4. Raster plots of the two multi-neuron chips during stimulus presentation. (a) and (b): raster plots of the multi-neuron chips respectively representing the X position and the Y position of the target. The red lines indicate the different phases of the stimulus, indicated by the colored arrows. The blue dots represent the input from the DVS. The time lag between the response of the neurons and the stimulation reflects the time constant of the synapses, which was in the order of hundreds of *ms*. From $t = 0s$ to $t = 1s$, we observe no activity because the target is stationary. Between $t = 1s$ and $t = 3s$ (gray arrow) the target moves horizontally to the right, and the DVS starts to stimulate the multi-neuron chips. At $t = 3s$, the target makes a 2s pause. Between $t = 5s$ and $t = 7s$ the target moves to the lower-right (represented by the green arrow) and finally stops. Even in the absence of sensory stimulation during $3s < t < 5s$ and $t > 7s$, we observe that the activity in the neurons persists. (c) and (d): raster plots of the two chips when presented with the incorrect stimulus (red arrow: target moving to the upper-right). During $3s < t < 5s$, the input stimulates a combination of populations which are not allowed to activate, given the relations between the two chips. As a result, the network resists the drift induced by the input. After $t > 7s$, the activity remains at its last valid location. In both cases, the activity occasionally moves in slight jumps. This is due to stronger attraction to the persistent activity states, and is caused essentially by mismatch in the transistors. The top-most neurons (124 to 128) are the inhibitory neurons. Mean activity of the inhibitory neurons activity $\cong 200\text{Hz}$, Mean activity of the excitatory neuron activity during persistent state $\cong 50\text{Hz}$.

[3] 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*, vol. 54, no. 1, pp. 48–59, Jan 2007, special Issue on Sensors.

[4] R. Douglas and K. Martin, "Neural circuits of the neocortex," *Annual Review of Neuroscience*, vol. 27, pp. 419–51, 2004.

[5] U. Rutishauser and R. Douglas, "State-dependent computation using coupled recurrent networks," *Neural Computation*, vol. 21, pp. 478–509, 2009.

[6] M. L. Minsky, *Computation: finite and infinite machines*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1967.

[7] P. Dayan and L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, 2001.

[8] R. Ben-Yishai, R. Lev Bar-Or, and H. Sompolinsky, "Theory of orientation tuning in visual cortex," *Proceedings of the National Academy of Sciences of the USA*, vol. 92, no. 9, pp. 3844–3848, April 1995.

[9] A. L. Yuille and N. M. Grzywacz, "A winner-take-all mechanism based on presynaptic inhibition feedback," *Neural Comput.*, vol. 1, no. 3, pp. 334–347, 1989.

[10] S. H. Strogatz, *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering (Studies in nonlinearity)*, 1st ed., ser. Studies in nonlinearity. Perseus Books Group, January 1994. [Online]. Available: <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0738204536>

[11] G. Indiveri, E. Chicca, and R. Douglas, "A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity," *IEEE Transactions on Neural Networks*, vol. 17, no. 1, pp. 211–221, Jan 2006.

[12] C. Bartolozzi, S. Mitra, and G. Indiveri, "An ultra low power current-mode filter for neuromorphic systems and biomedical signal processing," in *IEEE Proceedings on Biomedical Circuits and Systems (BioCAS06)*, 2006, pp. 130–133.

[13] J. Lazzaro, J. Wawrzynek, M. Mahowald, M. Sivilotti, and D. Gillespie, "Silicon auditory processors as computer peripherals," *IEEE Transactions on Neural Networks*, vol. 4, pp. 523–528, 1993.

[14] D. Fasnacht, A. Whatley, and G. Indiveri, "A serial communication infrastructure for multi-chip address event system," in *IEEE International Symposium on Circuits and Systems, ISCAS 2008*. IEEE, May 2008, pp. 648–651.

[15] P. Lichtsteiner and T. Delbrück, "A 64x64 AER logarithmic temporal derivative silicon retina," in *Research in Microelectronics and Electronics, 2005 PhD*, vol. 2, July 2005, pp. 202–205.

[16] A. D. Straw, "Vision egg: an open-source library for realtime visual stimulus generation," *Frontiers in Neuroinformatics*, vol. 2, 2008. [Online]. Available: frontiersin.org/neuroinformatics/paper/10.3389/neuro.11/004.2008/