

Form specifies function: Robust spike-based computation in analog VLSI without precise synaptic weights

Dylan R. Muir, Giacomo Indiveri and Rodney Douglas
Institute of Neuroinformatics
UNI - ETH Zurich
Zurich, Switzerland
Email: dylan@ini.phys.ethz.ch

Abstract—Recent studies demonstrated that sophisticated information processing can occur in spike-based computational systems that make use of synapses with only two states (potentiated or depressed). Here we present the hybrid software / hardware implementation of a model of the mammalian olfactory bulb using an analog VLSI device comprising an array of integrate and fire neuron with bistable synapses. Our implementation incorporates both software and hardware components, integrated using an asynchronous event-based spike representation. The model is able to perform highly selective simulated odor recognition, using induced synchronization within a population of neurons as the key to computation. The success of this scheme shows that the analog VLSI circuits used can perform sophisticated computation, taking advantage of the neuron dynamics and the topology of the network, without requiring precise analog synaptic weights.

I. INTRODUCTION

Networks of integrate and fire (I&F) neurons have been shown to exhibit a wide range of useful computational properties [1]. Moreover, recent studies have demonstrated that sophisticated information processing can occur in spike-based computational systems that make use of synapses with only two states (potentiated or depressed) [2]–[4]. These types of networks with bistable synapses are very well suited for analog VLSI implementation. In previous work we presented neuromorphic circuits for implementing massively parallel, distributed networks of I&F neurons [5], [6], and a VLSI device comprising both I&F neurons and plastic, bistable synapses, which can be used to implement multi-chip reconfigurable networks of spiking neurons [7].

Here we present an application of such a device in which we implement an olfactory bulb model, based on the model proposed in [8]. The implementation we propose uses both software and hardware components and demonstrates how it is possible to build hybrid software/hardware networks of I&F neurons with plastic, bistable synapses capable of performing complex tasks, such as highly selective odor recognition.

II. THE ANALOG VLSI DEVICE

The analog VLSI device was implemented using a standard AMS 0.8 μ m CMOS process. It comprises an array of 32 I&F neurons, a 2-D array of 32 \times 8 synaptic analog circuits, and

digital Input/Output circuits that allow the device to receive and transmit pulses (spikes) in input and transmit spikes in output using the *Address-Event Representation* (AER) [9], [10]. In AER, each spiking neuron is assigned an address. When a neuron generates a pulse its address is put on a digital bus. In this asynchronous protocol time representation is implicit, and analog signals are encoded by the inter-spike intervals between the addresses of their sending nodes. Input and output spikes (events) are transmitted as real-time asynchronous binary data streams that carry analog information in their temporal structure. Each event is represented by a binary word encoding the address of the sending node.

Inside the VLSI device each neuron is connected to 2 inhibitory and 6 excitatory synapses (see Fig. 1). The excitatory synapses contain circuits for implementing adaptive learning on short time scales, and circuits for driving the weight to one of two possible analog values on long time scales. The learning circuits implement a rule based on the spike-timing-dependent plasticity (STDP) mechanism [11]. In STDP the precise timing of spikes generated by the neurons determines the weight update. If a pre-synaptic spike arrives at the synaptic terminal before a post-synaptic spike is emitted, within a critical time window, the synaptic efficacy is increased. Conversely if the post-synaptic spike is emitted soon before the pre-synaptic one arrives, the synaptic efficacy is decreased. If on short time scales the synaptic weight is increased by the STDP circuits above a threshold, the bistability circuits will generate a constant weak positive current. In the absence of activity (and hence learning) this current will drive the weight toward its potentiated state; if the STDP circuits decrease the synaptic weight below the threshold, the bistability circuits will generate a negative current that in the absence of spiking activity will actively drive the weight toward the analog value encoding its depressed state.

The addresses of the spiking neurons, transmitted off-chip by the AER circuits, are received by a custom PCI-AER board [12]. In addition to logging, time-stamping and storing the spiking activity of the neurons, the PCI-AER board can send Address Events to arbitrary synapses on the chip. Hence spike trains with arbitrary statistics can be generated off-line

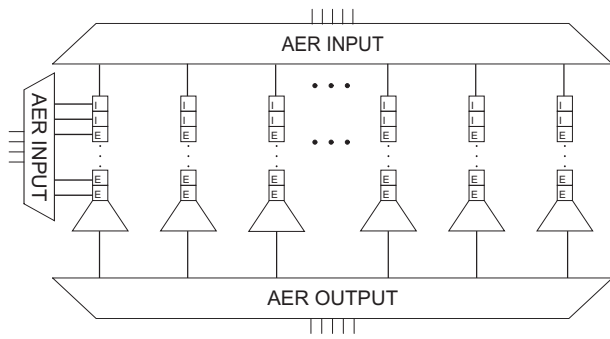


Fig. 1. Block diagram of the chip's architecture. Small trapezoids in the array represent integrate and fire neurons, while squares represent inhibitory and excitatory synapses.

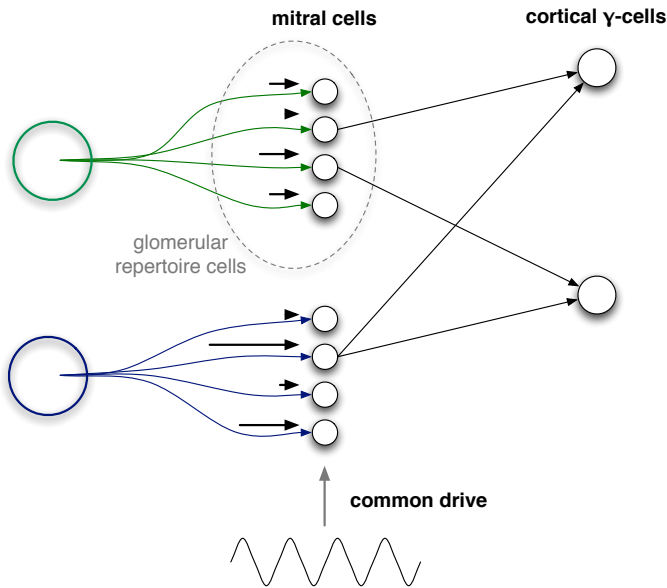


Fig. 2. Olfactory bulb model proposed by Brody and Hopfield. Individual neurons in the model are represented by circles. The two circles on the left represent glomerular cells. The activity of chemosensors (not shown) is proportional to the chemical components of an odor. Chemosensors sensitive to a particular component integrate their activity in a single glomerulus. Mitral cells, in the center of the figure, have three distinct input sources: odor-elicited activation from the glomeruli, a common sinusoidal drive, and a constant random bias current (bold arrows). The dashed oval indicates the repertoire of mitral cells receiving common input from a glomerulus.

on a workstation by a software model, mapped to arbitrary synapse addresses, depending on the desired network topology, and sent to the chip. The result of the network activity can be measured by monitoring the spikes generated by the neurons in the device.

III. THE OLFACTORY MODEL

Brody and Hopfield proposed a model of the mammalian olfactory bulb in a spike-based system, using induced synchronization within a population as the key to the computation performed by the model [8]. This model is shown in Fig. 2. Glomeruli collect and integrate the activity of chemosensors sensitive to a particular odor component. Each glomerular cell stimulates a repertoire of mitral cells with a common

activation. Each mitral cell in the glomerular repertoire is also stimulated by a randomly chosen constant bias current (bold arrows in the central part of Fig.2). In addition to the odor-elicited activations and the constant random currents, the mitral cell population receives a common sinusoidal current input. The effect of this common sinusoidal input is to induce approximate synchronization among the mitral cells that have similar *net* input currents. For a range of net input currents, the spiking activity of the population phase locks with the common sinusoidal input. For input currents outside this range, the mitral cells do not fire in synchrony.

The effect of odor presentation in the olfactory model is illustrated in Fig. 3. The components of an odor drive glomeruli with varying strengths. In the figure, the repertoire of mitral cells associated with each glomerulus is shown grouped together. Constant random bias currents injected into each mitral cell are shown as horizontal bars representing the magnitude of the bias current. The glomerular activations, when coupled with the random currents will cause a subset of the mitral cells to have a similar total activation. These mitral cells are highlighted in the figure, and the dashed line shows that the net activation for these cells is approximately equal. The highlighted mitral cells in Fig. 3 will fire in close synchrony. Connections are then manually selected from only those highlighted cells to a single gamma cell, which acts as a synchrony detector for this sub-population (see also Fig. 2), given that the membrane time constant of the gamma cell is shorter than the wavelength of the common sinusoidal input to the mitral cell population. The activations projected to a single gamma cell upon presentation of an odor are shown on the right of Fig.3. The top right box of Fig.3 shows the effect of presenting the target odor to the model. The mitral cells will fire in synchrony, inducing the gamma cell to fire. The lower right box of Fig.3 shows a representation of the input to the same gamma cell upon presentation of a non-target odor. As the same mitral sub-population will not fire in synchrony, the gamma cell is less likely to fire. The number of spikes produced by a gamma cell is a direct measure of its recognition of the odor presented to the model.

IV. THE HYBRID SOFTWARE/HARDWARE MODEL

We implemented this model in a hybrid software/hardware simulation. Odor generation, glomerular activity and mitral cell activity were contained in a software simulation of simple integrate-and-fire neurons. This portion of the system was adapted from Brody and Hopfield's matlab-based olfactory simulation, and incorporates the same model neurons outlined in their paper [8]. Both glomerular input and the common sinusoidal drive were injected as currents into the simulated mitral cells. Gamma cells were implemented by the silicon neurons on the analog VLSI device. Connections were made between the software portion and the hardware portion of the implementation using the custom PCI-AER board [12]. In our implementation, mitral cell spikes were generated in software and used to directly stimulate excitatory synapses on the analog VLSI neuron.

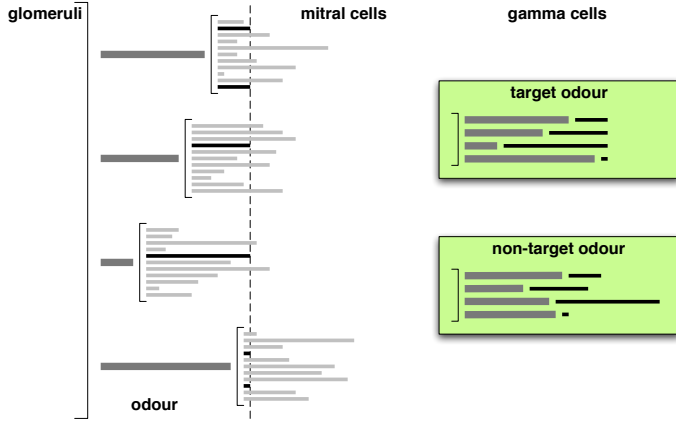


Fig. 3. Effect of simulated odor presentation in the olfactory model. Activations are represented by horizontal bars that encode the magnitude of the cell’s stimulating current. The activations of four glomeruli resulting from the response to an odor are shown on the left. Every mitral cell in a repertoire receives an identical odor-driven activation. In addition, each mitral cell receives a constant random bias current. The highlighted subset of the mitral cell population receives the same *net* activation, and will therefore fire in approximate synchrony. The highlighted cells are manually selected to project to a gamma cell. On the right, the net activation of the highlighted mitral cells is shown for the target odor (top box) and a non-target odor (bottom box).

The parameters of the analog VLSI neurons were set to imitate as closely as possible the time constants, threshold potential and other parameters of the gamma cells in Brody and Hopfield’s original simulations [8].

V. THE EXPERIMENT

In Hopfield and Brody’s network-repair extension to their olfactory model [4], the connections between synchronous subsets of mitral cells and the gamma cell are learned with a spike-timing based learning rule. Specifically, 200 synapses are potentiated from a field of 5,400 potential mitral cell connections. Due to the low numbers of excitatory synapses per neuron available on our aVLSI device we could not carry out an equivalent experiment, where the synapses of the gamma cell with appropriate connections (*i.e.* the synapses stimulated by the synchronous subset of mitral cells) potentiated and the synapses with inappropriate connections depress. Instead we manually selected the appropriate subset of mitral cells to connect to the aVLSI synapses of a gamma cell and verified that the STDP circuits maintained the weights in their potentiated state. Furthermore we evaluated the selectivity of the aVLSI gamma cell by gradually replacing the mitral cells projecting to the gamma cell synapses, from an initially appropriate synchronous subset to a randomly selected subset. According to the STDP learning rule implemented on the chip, the synapses receiving input from inappropriate mitral cells should decrease their weight to the depressed state, dramatically decreasing the sensitivity of the gamma cell to the target odor.

We generated random odor activity patterns in software and presented them to the simulated mitral cell population. We collected the spikes generated by the mitral cells, and selected a set of six mitral cells communally sensitive to the random odor. We mapped the output spikes from these mitral cells to the six excitatory synapses on an aVLSI neuron via the PCI-AER board. These connections caused the aVLSI neuron to behave as a gamma cell highly sensitive to the random odor. As a control, we created a non-selective gamma cell by connecting six random mitral cells to its six excitatory synapses. Upon a simulated odor presentation, the mitral cells produced spikes; the spikes were transmitted to the aVLSI gamma cells, and the spiking activity of the aVLSI gamma cells was recorded by the PCI-AER board. We repeated this procedure, each time progressively replacing individual appropriate mitral cell connections with random connections, until all six possible mitral cell connections were made randomly. For the engineered gamma cell, this had the effect of reducing its sensitivity. For the control (non-selective) gamma cell there was no significant reduction in sensitivity, as new random connections could cause the cell to become even more sensitive to the target odor. In all cases the net number of spikes received by a gamma cell was roughly equal.

VI. RESULTS

Figure 4 shows the results of the odor presentation trials. The response of the gamma cells is shown for progressive replacement of the mitral cell connections. The response marked with a double asterisk shows that the engineered gamma cell is highly selective for the target odor. This selectivity is statistically significant ($p < 0.001$, $n = 175$).

The number of spikes produced by a gamma cell during odor presentation is a direct measure of the degree to which it recognizes an odor. The selectivity of the model implementation can be measured by comparing the number of spikes produced in response to the target odor with the number of spikes produced in response to a random odor. In Fig. 4, this is equivalent to comparing the solid and dashed lines for zero inappropriate synaptic connections. For our implementation, this selectivity was 1.98 : 1.

When the network is deliberately damaged by degrading the appropriateness of the mitral-gamma cell synaptic connections, the model should continue to provide high selectivity to odors. The STDP-based learning rules implemented in each synapse cause inappropriate mitral-gamma cell synaptic connections to become depressed. This ensures that only synchronized mitral cell spikes contribute to the activity of a gamma cell. Incorporating STDP learning therefore depresses the activity of a gamma cell when presented with an odor close to the cell’s target odor. This effect can be seen by the sharpness of the selectivity between zero and one inappropriate synaptic connections for the manually engineered gamma cell (the solid line in Fig. 4). For our implementation, this selectivity was 1.78 : 1.

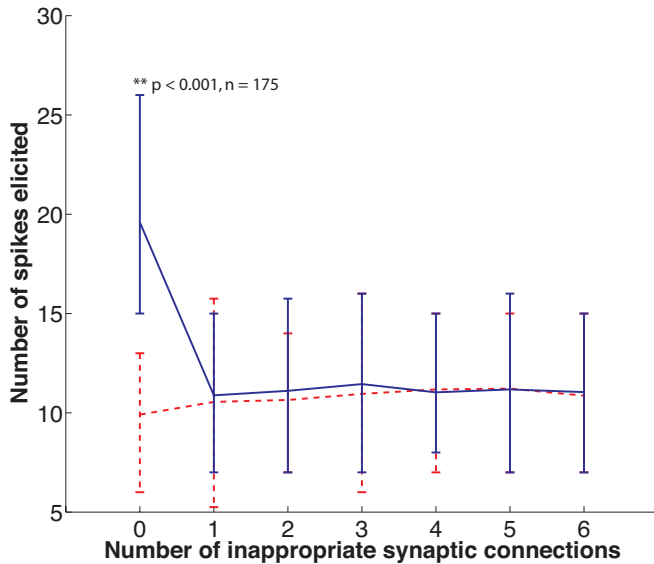


Fig. 4. Results of the odor presentation trials. Shown is the number of spikes elicited from a gamma cell on odor presentation, for an increasing number of inappropriate synaptic connections. The solid line represents a gamma cell for which the mitral cell connections were initially manually engineered to produce high odor sensitivity. The dashed line represents a gamma cell for which the initial connections were made randomly. Error bars show the upper and lower quartile of the response. The initial response of the engineered gamma cell is highly selective for the target odor. This response is indicated with a double asterisk, and is highly significant when compared to the response of the non-selective gamma cell ($p < 0.001$, $n = 175$).

VII. FURTHER WORK

The selectivity and function of the current implementation is limited mainly by the small number of synaptic inputs available to each gamma cell. The original model proposed in [4], where mitral-gamma cell connections are learned by an unsupervised learning rule, presupposes up to 200 mitral cell synaptic connections to each gamma cell. A greater number of synapses per gamma cell in our hybrid software/hardware model would allow greater odor selectivity, would lessen the effect of operating conditions on the aVLSI hardware and would allow more sophisticated network repair and unsupervised learning experiments to be performed. We are currently working on a new implementation of the model using a larger chip, recently fabricated, with 256 neurons and 32 synapses per neuron, for a total of 8,192 synaptic circuits. We will use this new implementation of the model to characterize its odor discrimination capabilities also for simultaneously presented odors.

VIII. CONCLUSION

We demonstrated that computation can occur in a hybrid software/hardware simulation using synaptic connections with only two states (potentiated or depressed). Using an analog VLSI device that incorporates bistable synapses with an STDP learning rules, and integrate and fire neurons, we implemented an olfactory model capable of performing robust odor discrimination. The model implementation was found to be highly

selective between target and non-target odors, in addition to maintaining this high selectivity while undergoing degradation of the network connectivity. The STDP-based learning rules implemented in the synapses of the analog VLSI device supported the selectivity of the model by limiting the effects of the network degradation.

The success of this scheme shows that sophisticated computation is possible in noisy networks of spiking neurons relying only on neuron dynamics and the structure of a spiking network. In this context, precise analog synaptic weights are not necessary to perform computation.

ACKNOWLEDGMENTS

This work was supported by the EU grant ALAVLSI (IST-2001-38099), by the ETH Zürich TH Project 02254, and by the John Crampton Travelling Fellowship. We would like to also acknowledge Vittorio Dante and Paolo Del Giudice, ISS Italy, for the original design of the PCI-AER board, and Adrian Whatley, Elisabetta Chicca, Matthias Oster and Gerd Dietrich for further development of PCI-AER board hardware and software components.

REFERENCES

- [1] W. Maass and C. M. Bishop, *Pulsed Neural Networks*. MIT Press, 1998.
- [2] S. Fusi, M. Annunziato, D. Badoni, A. Salamon, and D. Amit, "Spike-driven synaptic plasticity: theory, simulation, VLSI implementation," *Neural Computation*, vol. 12, pp. 2227-2258, 2000.
- [3] D. Floreano and J. Urzelai, "Neural morphogenesis, synaptic plasticity, and evolution," *Theory in Biosciences*, vol. 120, pp. 225-240, 2001.
- [4] J. J. Hopfield and C. D. Brody, "Learning rules and network repair in spike-timing-based computation networks," *Proc. Natl. Acad. Sci.*, vol. 101, no. 1, pp. 337-342, 2004.
- [5] G. Indiveri, "A low-power adaptive integrate-and-fire neuron circuit," in *Proc. IEEE International Symposium on Circuits and Systems*. IEEE, May 2003, pp. IV-820-IV823.
- [6] —, "Neuromorphic bistable VLSI synapses with spike-timing-dependent plasticity," in *Advances in Neural Information Processing Systems*, vol. 15. Cambridge, MA: MIT Press, December 2002, pp. 1091-1098.
- [7] G. Indiveri, E. Chicca, and R. Douglas, "A VLSI reconfigurable network of integrate-and-fire neurons with spike-based learning synapses," in *Proceedings of 12th European Symposium on Artificial Neural Networks (ESANN04)*, 2004, pp. 405-410.
- [8] C. D. Brody and J. J. Hopfield, "Simple networks for spike-timing-based computation, with application to olfactory processing," *Neuron*, vol. 37, pp. 843-852, 2003.
- [9] K. Boahen, "Communicating neuronal ensembles between neuromorphic chips," in *Neuromorphic Systems Engineering*, T. S. Lande, Ed. Norwell, MA: Kluwer Academic, 1998, pp. 229-259.
- [10] S. R. Deiss, R. J. Douglas, and A. M. Whatley, "A pulse-coded communications infrastructure for neuromorphic systems," in *Pulsed Neural Networks*, W. Maass and C. M. Bishop, Eds. MIT Press, 1998, ch. 6, pp. 157-178.
- [11] L. F. Abbott and S. Song, "Asymmetric hebbian learning, spike timing and neural response variability," in *Advances in Neural Information Processing Systems*, vol. 11, 1998, pp. 69-75.
- [12] V. Dante and P. Del Giudice, "The PCI-AER interface board," in *2001 Telluride Workshop on Neuromorphic Engineering Report*, A. Cohen, R. Douglas, T. Horiuchi, G. Indiveri, C. Koch, T. Sejnowski, and S. Shamma, Eds., 2001, pp. 99-103, <http://www.ini.unizh.ch/telluride/previous/report01.pdf>.