ine Institute of Neuromorphic Engineering

THE NEUROMORPHIC ENGINEER
A PUBLICATION OF INE-WEB.ORG

10.2417/1200904.1636

Spike-based synaptic plasticity and classification on VLSI

Srinjoy Mitra and Giacomo Indiveri

We describe a VLSI system that implements a bioplausible spike-based learning algorithm and is capable of robust classification of binary patterns, even when they are highly correlated.

The important role of activity-dependent modifications of synaptic strength in learning and memory formation is well accepted in the neuroscience community.[1] The process of memory formation in neural networks is usually the result of a training procedure during which the synaptic strengths between neurons are modified according to a learning rule. During training, the network has to create, modify, and preserve memories of the representations of the learned classes with assistance from a supervisory input. During testing, the network—when presented with patterns belonging to these classes—should be able to identify them without the supervisor.

Here we describe a VLSI system, based on a network of integrate-and-fire (I&F) neurons and plastic synapses,[2] that can learn to classify complex patterns of mean firing rates into binary classes. The learning rule is based on a recently proposed model of stochastic spike-driven synaptic plasticity that can encode patterns of mean firing rates, and captures the rich phenomenology observed in neurophysiological experiments.[3] The theory requires the synapses show long-term plasticity that is bistable in nature: they rest in either a potentiated or depressed state after memory formation has taken place. This learning rule is ideally suited to large-scale silicon implementations, as long-term storage and retrieval of binary values is easily supported in standard VLSI technology.

In Figure 1 we show a simplified block diagram of a neuron along with its synapses.[4] The VLSI chip consists of 16 similar neurons each with 60 plastic and 4 non-plastic synapses. The synapses receive spike-train input from a PC or a spike-based sensor chip via an asynchronous event-driven protocol (we use AER, address event representation). Using the same protocol, the neuron sends its spikes off-chip to a PC for data logging and further processing. For the classification experiment, the plastic synapses were presented with a spatial pattern of

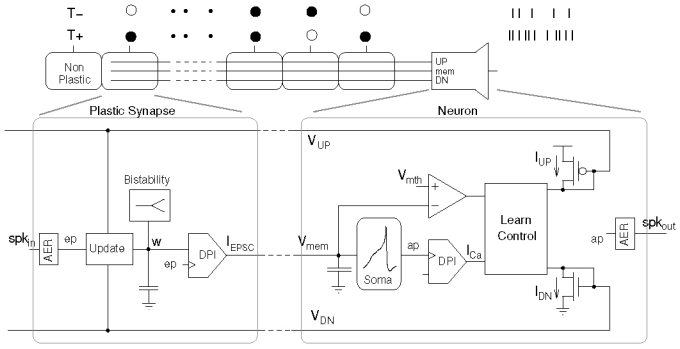

Figure 1. *Schematic diagram of a neuron and its synapses with details of its basic functional blocks. The plastic synapses are presented with low/high (binary) patterns of mean firing rates of either 2Hz (black circles) or 30Hz (white circles). The non-plastic synapses have a fixed weight and are stimulated with a supervisor signal that corresponds to a spike train with either very-high or very-low mean firing rate ($T+$ or $T-$).*

stimuli consisting of Poisson-distributed spike trains with either low (2Hz) or high (30Hz) mean firing rate. During the training phase, an additional supervisory input—in the form of a Poisson-distributed spike train with low ($T-$) or high ($T+$) spike rates—is presented to the non-plastic synapse. The high/low value of the mean firing rate is determined by the class to which the randomly-generated binary input pattern belongs ($C-$ or $C+$). The Poisson nature of the spike-trains is required to comply with the model requirements.

During training, the neuron responds with mean firing rates that reflect the input received from the supervisory signal and (in part) from the plastic synapses. Depending on these mean firing rates, the neuron circuit activates or deactivates two feedback signals (UP and DN) conveyed in parallel to all plastic synapses. The weight of each plastic synapse (w) is updated at the arrival of a corresponding pre-synaptic spike (spk_{in}) with an upward or downward jump that depends on the value of the UP/DN feedback signals from the neuron. The *bistability* block is responsible for the long-term dynamics of the synapse and the DPI

Continued on next page

(diff-pair integrator) circuit[5] integrates the pre-synaptic spikes into a synaptic current (I_{EPSC}) that is sourced into the neuron's membrane capacitance C_{mem}. All synaptic currents are summed spatially in parallel onto the neuron's membrane capacitance. The resulting output spikes (*ap* in Figure 1) are further integrated by another instance of the DPI circuit into the current I_{Ca}. They are used in the *Learn Control* block to activate the feedback signals UP (if V_{mem} is above a set threshold V_{mth} and I_{Ca} is within set bounds) or DN (if V_{mem} is below V_{mth} and I_{Ca} is within set bounds).

The pre-synaptic signals labeled *ep*, at the input of each plastic synapse in Figure 1, are voltage pulses produced by *AER* interface circuits. During pre-synaptic activity, the weights undergo instantaneous up/down jumps that depend on the status of the UP and DN signals produced by neuron's learning circuit. In the absence of pre-synaptic spiking activity, the bistability circuits drive the weights to the high (potentiated) state if they are above a set threshold, or to the low (de-potentiated) state if they are below it. In the case of consolidated transitions from low to high, the synapse is said to have performed long term potentiation (LTP). Similarly, in case of high to low transitions, the synapse undergoes long term depression (LTD).

Given the Poisson nature of the input spike trains, the plastic synapses undergo stochastic LTP/LTD transitions. However, as the synapses are trained with multiple repetitions of the same mean pre-synaptic and post-synaptic firing rates, the neuron eventually learns to respond correctly to the input patterns (*i.e.* with the expected high or low mean firing rate). After training, the LTP/LTD transitions are consolidated by the bi-stability circuits. In accordance with the theoretical model,[3] the probabilities of LTP and LTD are highly dependent on the mean post-synaptic firing frequency (ν_{post}). We verified this by stimulating 60 synapses in a single neuron with controlled pre- and post-synaptic spike-frequencies. We increased the neuron's output firing rate ν_{post} by driving it via a non-plastic synapse and used, as pre-synaptic inputs to the plastic synapse, Poisson distributed spike trains of 60Hz. In Figure 2 we show the LTD (top row) and LTP (bottom row) transitions obtained by setting the synapse initial state to the potentiated or depotentiated state respectively, before starting the stimulation. We repeated the same experiment 20 times with identical pre and post-synaptic mean frequencies but different instantiations of the Poisson statistics. As evident from the plots, the synaptic transitions are random in nature but with a probability that depends on ν_{post}. The figure inset shows the average LTP and LTD transitions for all sixty synapses across all trials. These curves are a good fit with the theoretical predictions made in the modeling work.[3]

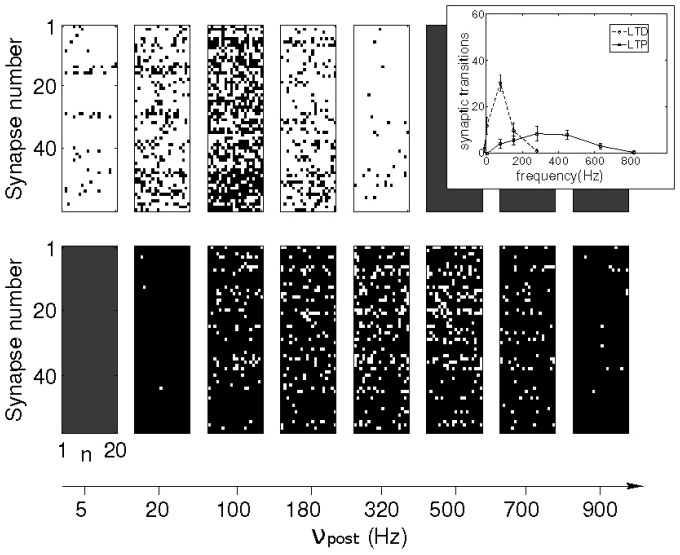

Figure 2. *Experimental results demonstrating the stochastic transition probability in silicon synapses. The black dots in the top row represent a long-term-depression (LTD) transition while the white dots in the bottom row represents a long-term-potentiation (LTP) transition. Gray boxes indicate no data for that particular post-synaptic frequency (ν_{post}). The figure inset shows the average LTD and LTP transitions across trials as a function of ν_{post}.*

To quantify the learning and classification properties of the VLSI network, we carried out an extensive set of experiments and performed statistical analysis of the results. Neurons were trained with a large number of randomly generated binary patterns (as depicted in Figure 1) that were randomly assigned to class $C+$ or class $C-$. During the testing phase, we measured the average post-synaptic frequencies and used these results for generating standard ROC (receiver operating characteristics)[6] plots and to quantify the binary neuronal classifier's performance. Higher values of the 'area under the ROC curve' (AUC) measures indicate better classifier performance. Our results[4] produced AUC values above 0.85 for binary classification of sets of 2, 4, 6 and 8 random patterns with 60 synapses, indicating excellent classification performance.[6]

One key aspect of the learning rule implemented on our chip is its ability to robustly classify patterns into binary classes, even if they are highly correlated. To verify this, we presented the chip with multiple sets of binary patterns and increased the amount of (spatial) correlation among them. As predicted by theory, higher numbers of patterns to be simultaneously classified result

Continued on next page

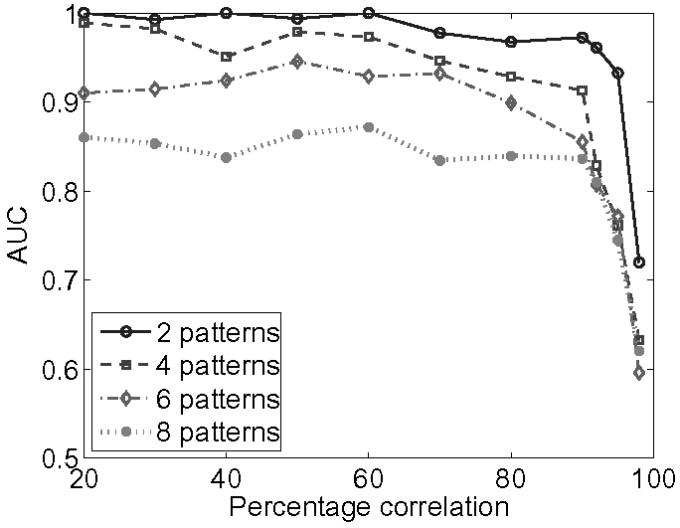

Figure 3. *Area under the curve (*AUC*) computed from the ROC analysis for different sets of patterns (2 to 8), as a function of the percentage of correlation among the patterns.*

in lower AUC values. This is evidenced in the four sets of curves plotted in Figure 3 (*e.g.* for a fixed value of percentage correlation). However, given a fixed number of patterns, the network is much less affected by the increased amount of correlation among the patterns, and degrades smoothly when the correlations render the patterns almost indistinguishable.[4]

In summary, we presented results from a mixed-signal VLSI chip capable of performing spike-based synaptic plasticity leading to memory formation and classification of binary patterns. We carried out extensive testing to verify the chip's robustness in its classification behavior, and demonstrated its performance even with strongly correlated input patterns. To our knowledge, the performance achieved with this system has not yet been reported for any other spike-based learning VLSI chips. The device proposed is very attractive for neuromorphic engineers, as it could be efficiently exploited for a wide range of sensory-motor applications including on-line learning in autonomous robotics, real-time spike-based computational module in brain-machine interfaces, and so on.

This work was supported by Swiss National Science Foundation grant No. PP00A1 06556, ETH grant No. TH0201 7404, and EU grants ALAVLSI (IST-2001-38099) and DAISY (FP6-2005-015803).

Author Information

Srinjoy Mitra
Computational Sensory-Motor Lab
The Johns Hopkins University
Baltimore, MD

Srinjoy Mitra is a post-doctoral researcher. He received his PhD from the ETH Zurich in 2008 working on spike-based computation in VLSI. His current interests are neuromorphic circuits and analog VLSI design for bio-engineering applications.

Giacomo Indiveri
Institite of Neuroinformatics (INI)
University of Zurich and ETH Zurich
Zurich, Switzerland

Giacomo Indiveri is a Senior Lecturer at INI. He obtained his PhD in electrical engineering from the University of Genoa, Italy and his research interests are in the domains of neuromorphic circuits, selective attention systems, winner-take-all models, and VLSI spike-based computational systems.

References

1. P. Dayan and L. F. Abbott, **Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems**, MIT Press, 2001.
2. G. Indiveri and S. Fusi, *Spike-based learning in VLSI networks of integrate-and-fire neurons*, **Proc. IEEE Int'l Symp. on Circuits and Systems (ISCAS)**, pp. 3371–3374, 2007.
3. J. Brader, W. Senn, and S. Fusi, *Learning real world stimuli in a neural network with spike-driven synaptic dynamics*, **Neural Computation 19**, pp. 2881–2912, 2007.
4. S. Mitra, S. Fusi, and G. Indiveri, *Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI*, **IEEE Trans. on Biomedical Circuits and Systems 3** (1), pp. 32–42, Feb. 2009.
5. C. Bartolozzi and G. Indiveri, *Synaptic dynamics in analog VLSI*, **Neural Computation 19** (10), pp. 2581–2603, Oct 2007.
6. T. Fawcett, *An introduction to ROC analysis*, **Pattern Recognition Lett. 26**, pp. 861–874, 2006.