A hybrid analog/digital Spike-Timing Dependent Plasticity learning circuit for neuromorphic VLSI multi-neuron architectures

Hesham Mostafa, Federico Corradi, Fabio Stefanini, and Giacomo Indiveri
Institute of Neuroinformatics, University of Zürich and ETH Zürich, Switzerland
Email: {hesham,federico,fabio.stefanini,giacomo}@ini.uzh.ch

***Abstract*—To endow large scale VLSI networks of spiking neurons with learning abilities it is important to develop compact and low power circuits that implement synaptic plasticity mechanisms. In this paper we present an analog/digital Spike-Timing Dependent Plasticity (STDP) circuit that changes its internal state in a continuous analog way on short biologically plausible time scales and drives its weight to one of two possible bi-stable states on long time scales. We highlight the differences and improvements over previously proposed circuits and demonstrate the performance of the new circuit using data measured from a chip fabricated using a standard 180 nm CMOS process. Finally we discuss the use of stochastic learning methods that can best exploit the properties of this circuit for implementing robust machine-learning algorithms.**

I. Introduction

In recent years, there has been a proliferation of custom VLSI implementations of computing architectures based on spiking neural networks [1]–[6]. These architectures compute in a distributed fashion that is massively parallel and fault tolerant and they can potentially have very low power consumption. Therefore, they are being explored as possible alternative solutions to conventional digital computing approaches. The "programming model" for these networks, however, is radically different from the conventional von Neumann sequence of instructions model. The knowledge, or the program, in a neural network is encoded in the network architecture, and in the weights of the connections between the neurons. New theories and methods are being developed for understanding how to synthesize pre-specified computing algorithms on these hardware systems through defining appropriate network architectures and connectivity patterns [7], [8]. However, to endow these architectures with adaptation and learning mechanisms that change their synaptic weights on-line (e.g., to realize the intended functionality autonomously) it is necessary to develop appropriate compact and low-power plasticity circuits that are compatible with the VLSI architectures being developed. Furthermore, to allow these hardware neuromorphic systems to interact with the user and the environment (e.g., in the field of robotics or neuroprosthetics), it is important that these circuits can exhibit dynamics with biologically plausible time constants.

In this paper, we describe a hybrid analog/digital weight-update circuit that can be biased to operate with such time constants. The circuit can be used in a wide variety of spike-based learning protocols and is optimized for minimum size and power-consumption. As the circuit updates its internal state variable when it is stimulated by input digital pulses, it is ideally suited to implement Spike-Timing Dependent Plasticity (STDP) learning algorithms. The synapse circuit updates its state in an analog way, continuously over time, but it settles to one of two possible states on long-time scales. This enables the design of an extremely compact and low-power circuit that is robust to noise in both the circuit and the input signals.

We describe the circuit in the next section. In Section III we demonstrate its properties with experimental results measured from a recent 180 nm CMOS design. In Section IV we present our conclusions and discuss the advantages of binary synapses and show how they can be used for learning uncorrelated patterns and for implementing classical neural network learning algorithms.

II. The learning synapse circuit

The learning synapse circuit was integrated in a multi-neuron prototype chip, fabricated using a standard 180 nm CMOS process, comprising analog (adaptive exponential) integrate-and-fire neuron circuits, linear and non-linear synapse dynamics circuits, and digital asynchronous spike-based communication and processing circuits.

Figure 1 shows the schematic diagram of the learning circuit. On the arrival of an input pulse spike_pre, representing a pre-synaptic spike, the circuit produces an output synaptic current I_syn that can be summed to the output currents of all other synapses, possibly fed through a linear integrator circuit that models synaptic dynamics, and conveyed to the post-synaptic integrate and fire neuron circuit. Multiple instances of these circuits can in turn be embedded in a network of spiking neurons for carrying out neural processing tasks. The amplitude of each pre-synaptic current I_syn depends on the corresponding circuit's synaptic weight. In the circuit of Fig. 1, the synaptic weight is encoded by the voltage Vc: if, during a spike, Vc<wthr! then the output current I_syn is zero (the synapse is in a Long-Term Depressed –LTD– state); conversely if Vc>wthr!, then the output current I_syn is equal to the Operational Transconductance Amplifier (OTA) bias current (the synapse is in its Long-Term Potentiated –LTP– state). During (and only during) the spike, the OTA bias current is set by the wscale! analog bias. In the absence of spikes,

978-1-4799-3432-4/14/$31.00 ©2014 IEEE

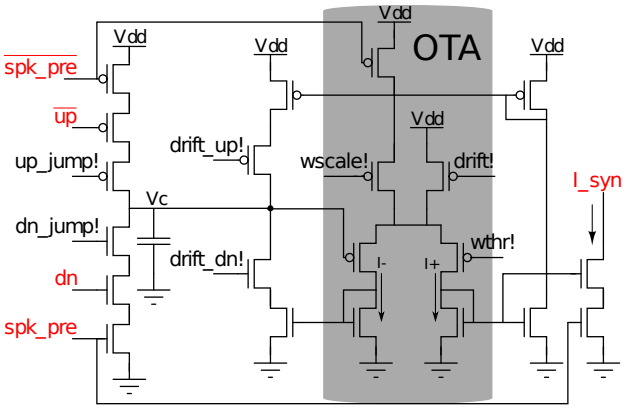

Fig. 1: The low-power analog/digital synapse circuit. Pre-synaptic input spikes (spk_pre) trigger the weight update that can be positive or negative depending on the digital up and dn signals. Red text indicates port names. Net Names ending in "!" indicate externally supplied biases.

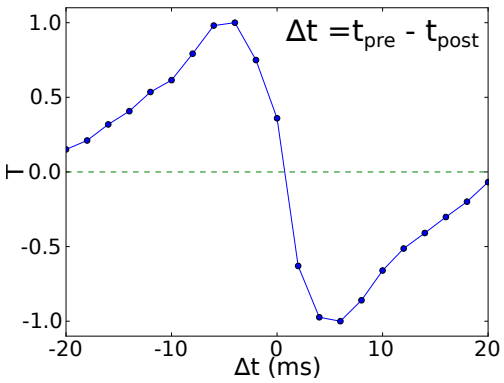

Fig. 2: Spike-timing dependent transition tendency as a function of Δt between pre- and postsynaptic spikes.

the OTA is biased by a typically much smaller current, set by the drift! bias voltage, which slowly drives the Vc voltage towards the Vdd power rail if Vc>wthr! and towards Gnd if Vc<wthr!. The drift towards one of the two stable states can be extremely slow, and the power consumption can consequently be very small, because the OTA is configured in a positive-feedback configuration and it can operate correctly even with extremely small biases (e.g., using only leakage currents). The analog drift_up! and drift_dn! signals can be used to control the slopes of the up and down drifts independently. Learning takes place by updating the voltage Vc with each pre-synaptic spike: during a spike, the charge on the Vc capacitor can be either increased, decreased, or left unchanged, depending on the values of the control signals up and dn. The analog biases up_jump! and dn_jump! determine the amount of charge being dumped or sourced from the capacitor (and therefore the amplitude of the Vc change).

The up and dn control signals can be shared among all synapses afferent to a neuron. Therefore only one circuit producing these signals needs to be instantiated (e.g., next to the neuron circuit). These signals can be generated based on post-synaptic spikes or based on the membrane potential of the postsynaptic neuron. We show in section III that generating the up and dn signals based only on the membrane potential is enough to reproduce the classical STDP learning algorithm [9]. However, more elaborate and powerful models can be implemented by using circuits that take into account additional postsynaptic variables [10]. For example, in the prototype chip that includes the synapse described, the circuits producing the up and dn signals are based on a spike-based learning algorithm originally described in [11], and successfully applied to pattern classification in neuromorphic hardware [2], [12].

A. Comparison to prior work

There are two circuits that have been presented in the past, which are closely related to the learning circuit we propose. The synaptic circuit described in [13] has the same functionality of the synapse of Fig. 1: it generates either a high or low current, with every pre-synaptic spike, depending on the state of the internal synaptic weight variable, and drives the internal variable to a high or low state on longer-time scales. However this circuit uses two distinct OTAs (one for the comparison, and one for the drift). In addition to being larger, the OTA performing the comparison between the internal variable and a threshold is constantly active and dissipates power continuously (as opposed to the circuit of Fig. 1 which activates this feature only during spikes, when it is needed).

The other related learning circuit is the one that was originally described in [14]. In this prior implementation the output synaptic current was not a threshold function of the internal variable Vc, but was directly proportional to the exponential of Vc (by means of a sub-threshold nFET). This was a more compact and low-power circuit, but with an extremely high sensitivity to weight fluctuations. We carried out experiments to highlight the differences between the two approaches (shown in Section III), and point out the advantages of the new solution.

III. Experimental Results

To show how the learning circuit of Fig. 1 is compatible with STDP learning rules, we generated a Poisson distributed teacher spike train with a mean firing rate of 200 Hz sent via a fixed synapse. This teacher spike train causes the post-synaptic neuron to fire stochastically with a mean rate of 15 Hz. At the same time, we stimulated the learning synapse with a spike (pre-synaptic) shifted in time with respect to the post-synaptic spike and we observed the behavior of the circuit for different $\Delta t = t_{pre} - t_{post}$ spike times. Following the voltage-dependent weight update rule described in [11], our circuit exhibited higher probabilities of weight increase (Vc receiving the current generated through up_jump!) when pre-synaptic spikes preceded post-synaptic spikes. Similarly, the probability of weight decrease was larger for cases when post-synaptic spikes preceded pre-synaptic spikes. To quantify these

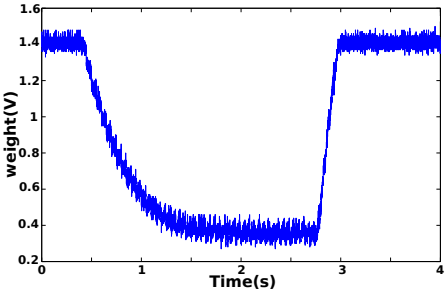

Fig. 3: Slow drift of the synaptic weight. The biases were changed at 0.4s and at 2.7s to trigger down and up drifts respectively.

effects, we computed the "tendency" function T, as defined in [11]:

$$T = \left(\frac{p_{up}}{p_{dn} + p_{up}} - \frac{1}{2} \right) \cdot max(p_{up}, p_{dn}) \qquad (1)$$

where p_{up} and p_{dn} are the probabilities of up and down weight changes for a given timing difference Δt. The plot of this tendency T as a function of Δt is shown in Fig. 2. The figure shows that the learning circuits can reproduce the STDP behavior.

In the next experiment, we measure the circuit's ability to realize slow temporal dynamics. This feature allows the network to process spike trains with low mean rates and biologically plausible temporal dynamics. Figure 3 illustrates the slow dynamics of the synaptic weight when it is drifting in the absence of presynaptic spikes. The drift! and the drift_up! bias voltages of Fig. 1 were both set to 1.8V, drift_dn! was set to 0.02V. wthr! was switched from 0V to 1.8V at 0.4s and back to 0V at 2.7s. Using such slow drift rate, weight updates can be sensitive to temporal correlations at the 1s scale in the presynaptic spike train.

In the experiment of Fig. 4, we compare the effect of learning in the synapse proposed in this work with the synapse previously proposed in [14]. Measurements were carried out on two separate chips that implement the two versions of the circuit. A neuron in each chip was stimulated via the plastic synapse. We set the synaptic weights to their high states in both circuits, set the biases in a way to induce down jumps in the Vc nodes, and stimulated the synapses with regular 200 Hz pre-synaptic spike trains. The stimulation starts at $t = 0.05$ s. The top plot shows the measurements from the prior implementation: even a very small reduction in the internal synaptic weight variable drastically alters the current injected into the neuron. As soon as the weight drops slightly, the synaptic output current becomes so small that it is not able to drive the neuron anymore. In contrast to this behavior, the bottom plot shows that changes in the internal variable Vc (labeled as W in the figure legend) have an effect only if they make it cross the threshold bias wthr! of Fig. 1 (labeled as W_{thr} in Fig. 4). Therefore this synapse is robust to small fluctuations in its synaptic weight internal variable (e.g., induced by noisy

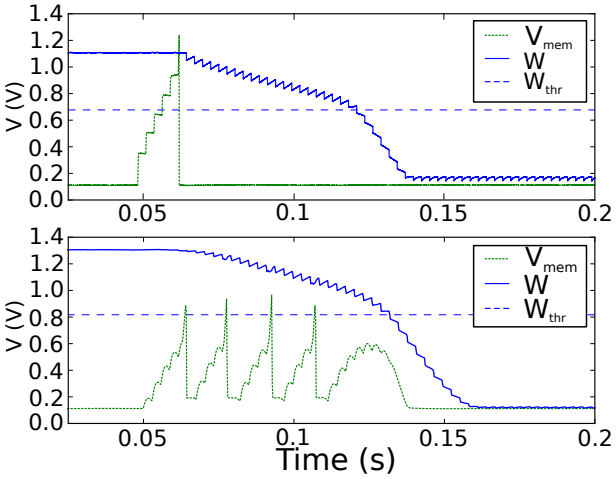

Fig. 4: Effect of fluctuations in the synaptic weight on the response of the postsynaptic neuron. The top plot shows measurements from the prior learning circuit implementation. The bottom plot shows measurements from the current implementation. Synaptic weights are actively driven to a low state. When the weights W are above the threshold W_{thr}, the down jumps act against a slow drift towards the high state. As soon as the weights cross W_{thr}, the drift acts in the same direction of the down jumps.

input spikes), and consolidates the changes that it is required to learn only if there are enough consistent signals that drive it to make a transition (note the drift component in the W traces of Fig. 4). We investigated the effect of the presynaptic weight on the neuron firing rate. The firing rate of a neuron driven by a regular spike train through the plastic synapse only changes when the synaptic weight crosses the bi-stability threshold wthr!.

IV. Discussion and Conclusions

In this work, we presented novel circuits for implementing bi-stable synapses, characterized their properties with experimental measurements, and showed the improvements made over previous implementations. The synapse proposed is extremely compact and low power, at the cost of having only two stable states as possible synaptic weights. The choice of implementing binary synapses with analog circuits may seem strange. However, it has been demonstrated that, despite the limited resolution, binary synapses are sufficient to learn random uncorrelated patterns of spike-trains, as in the perceptron case: perceptrons that use binary, stochastic synapses have been shown to work in software and hardware tasks in the past [2], [11], [13], [15], [16], and represent a good compromise between implementation complexity and functionality. In order to allow the realization of online linear classifiers with these synapses, it is important to use stochastic learning paradigms [15] and realize circuits where the synaptic efficacy is affected by the neuron only when the synaptic state crosses the bi-stability threshold [13]. The bi-stability

threshold bias (wthr! of Fig. 1) can further be used to change the balance between potentiation or depression without affecting the dynamics of the post-synaptic neuron, for example to realize reward-modulated learning or to compensate for circuit inhomogeneities due to mismatch.

The variability embedded in input (or natural) spike-trains can be effectively used as a source of stochasticity for implementing stochastic learning without requiring additional sources of noise or random number generators. Variability or noise in the input will not degrade performance: only when a sufficient number of spikes consistently impose potentiation or depression on a synapse, its internal state will counteract the drift current and undergo an LTP or LTD transition respectively (e.g., see LTD transition in Fig. 4).

Networks of silicon neurons that have these types of synapses and that use stochastic learning can be used effectively as an ensemble of linear classifiers trained with supervised learning to solve non-linear classification tasks. In this way, synapses corresponding to the same inputs connected to different post-synaptic neurons will undergo independent stochastic processes. Hence, even though the neurons in the ensemble of classifiers are trained with the same data, the resulting classifiers will have slightly different responses. By aggregating the responses of these neurons, a single, optimal, classifier can be realized in a single-layer neural network using the bi-stable synapses and the circuits we proposed [15]. This strategy is analogous to bootstrap aggregation methods in machine learning where multiple classifiers are independently trained and their responses combined [17]. Furthermore, by allowing only a subset of synapses to make synaptic transitions upon the presentation of the patterns to be learned, the system "samples" from a distribution of independent classifiers, with the result that the learned perceptron with bi-stable synapses will have better performance than a deterministic one [18].

ACKNOWLEDGMENT

This work was supported by the European Commission through the "neuroP" Project, ERC-2010-StG-257219-neuroP.

REFERENCES

[1] J. Arthur, P. Merolla, F. Akopyan, R. Alvarez, A. Cassidy, A. Chandra, S. Esser, N. Imam, W. Risk, D. Rubin, R. Manohar, and D. Modha, "Building block of a programmable neuromorphic substrate: A digital neurosynaptic core," in *International Joint Conference on Neural Networks, IJCNN 2012*. IEEE, Jun 2012, pp. 1946–1953.

[2] M. Giulioni, P. Camilleri, M. Mattia, V. Dante, J. Braun, and P. D. Giudice, "Robust working memory in an asynchronously spiking neural network realized in neuromorphic VLSI," *Frontiers in Neuroscience*, vol. 5, no. 149, 2012.

[3] S. Choudhary, S. Sloan, S. Fok, A. Neckar, E. Trautmann, P. Gao, T. Stewart, C. Eliasmith, and K. Boahen, "Silicon neurons that compute," in *Artificial Neural Networks and Machine Learning – ICANN 2012*, ser. Lecture Notes in Computer Science, A. Villa, W. Duch, P. Érdi, F. Masulli, and G. Palm, Eds. Springer Berlin / Heidelberg, 2012, vol. 7552, pp. 121–128.

[4] T. Yu, J. Park, S. Joshi, C. Maier, and G. Cauwenberghs, "65k-neuron integrate-and-fire array transceiver with address-event reconfigurable synaptic routing," in *Biomedical Circuits and Systems Conference, (BioCAS), 2012*. IEEE, Nov. 2012, pp. 21–24.

[5] D. Brüderle, M. Petrovici, B. Vogginger, T. Pfeil, S. Millner, A. Grübl, K. Wendt, E. Müller, M.-O. Schwartz, S. Jeltsch, J. Fieres, P. Müller, O. Breitwieser, L. Muller, A. Davison, J. Kremkow, M. Lundqvist, J. Partzsch, S. Scholze, L. Zühl, C. Mayr, A. Destexhe, M. Diesmann, T. Potjans, A. Lansner, R. Schüffny, J. Schemmel, and K. Meier, "A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems," *Biological cybernetics*, vol. 104, no. 4, pp. 263–296, 2011.

[6] G. Indiveri, B. Linares-Barranco, T. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Häfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen, "Neuromorphic silicon neuron circuits," *Frontiers in Neuroscience*, vol. 5, pp. 1–23, 2011.

[7] E. Neftci, J. Binas, U. Rutishauser, E. Chicca, G. Indiveri, and R. Douglas, "Synthesizing cognition in neuromorphic electronic systems," *Proceedings of the National Academy of Sciences*, vol. 110, no. 37, pp. E3468–E3476, 2013.

[8] C. Eliasmith, T. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen, "A large-scale model of the functioning brain," *Science*, vol. 338, no. 6111, pp. 1202–1205, 2012.

[9] L. Abbott and S. Nelson, "Synaptic plasticity: taming the beast," *Nature Neuroscience*, vol. 3, pp. 1178–1183, November 2000.

[10] J. Lisman and N. Spruston, "Postsynaptic depolarization requirements for ltp and ltd: a critique of spike timing-dependent plasticity," *Nature Neuroscience*, vol. 8, no. 7, pp. 839–841, Jul 2005.

[11] J. Brader, W. Senn, and S. Fusi, "Learning real world stimuli in a neural network with spike-driven synaptic dynamics," *Neural Computation*, vol. 19, pp. 2881–2912, 2007.

[12] S. Mitra, G. Indiveri, and S. Fusi, "Learning to classify complex patterns using a VLSI network of spiking neurons," in *Advances in Neural Information Processing Systems 20*, J. Platt, D. Koller, Y. Singer, and S. Roweis, Eds. Cambridge (MA): MIT Press, 2008, pp. 1009–1016.

[13] D. Badoni, M. Giulioni, V. Dante, and P. Del Giudice, "An aVLSI recurrent network of spiking neurons with reconfigurable and plastic synapses," in *International Symposium on Circuits and Systems, (ISCAS), 2006*. IEEE, May 2006, pp. 1227–1230.

[14] G. Indiveri and S. Fusi, "Spike-based learning in VLSI networks of integrate-and-fire neurons," in *International Symposium on Circuits and Systems, (ISCAS), 2007*. IEEE, 2007, pp. 3371–3374.

[15] W. Senn and S. Fusi, "Learning Only When Necessary: Better Memories of Correlated Patterns in Networks with Bounded Synapses," *Neural Computation*, vol. 17, no. 10, pp. 2106–2138, 2005.

[16] S. Mitra, S. Fusi, and G. Indiveri, "Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 3, no. 1, pp. 32–42, Feb. 2009.

[17] L. Breiman, "Bagging predictors," *Machine Learning*, vol. 24, pp. 123–140, 1996.

[18] G. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Improving neural networks by preventing co-adaptation of feature detectors," 2012, submitted on 03 Jul. 2012.