Implementation of a spike-based perceptron learning rule using TiO$_{2-x}$ memristors

Hesham Mostafa[1]*, Ali Khiat[2], Alexander Serb[2], Christian G. Mayr[1], Giacomo Indiveri[1] and Themis Prodromakis[2]

[1] Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland, [2] Nanoelectronics and Nanotechnology Research Group, School of Electronics and Computer Science, University of Southampton, UK

Synaptic plasticity plays a crucial role in allowing neural networks to learn and adapt to various input environments. Neuromorphic systems need to implement plastic synapses to obtain basic "cognitive" capabilities such as learning. One promising and scalable approach for implementing neuromorphic synapses is to use nano-scale memristors as synaptic elements. In this paper we propose a hybrid CMOS-memristor system comprising CMOS neurons interconnected through TiO$_{2-x}$ memristors, and spike-based learning circuits that modulate the conductance of the memristive synapse elements according to a spike-based Perceptron plasticity rule. We highlight a number of advantages for using this spike-based plasticity rule as compared to other forms of spike timing dependent plasticity (STDP) rules. We provide experimental proof-of-concept results with two silicon neurons connected through a memristive synapse that show how the CMOS plasticity circuits can induce stable changes in memristor conductances, giving rise to increased synaptic strength after a potentiation episode and to decreased strength after a depression episode.

Keywords: synaptic plasticity, silicon neurons, memristors, neuromorphic architectures, learning, perceptron

1. Introduction

Biological networks provide a tantalizing proof of the existence of a physically implementable computing architecture that is distributed, fault-tolerant, adaptive, and that outperforms conventional architectures in many important problems such as visual processing and motor control. This has motivated the development of various neuromorphic computing systems whose architectures reflect the general organizational principles of nervous systems in an effort to partially reproduce the immense efficiency advantage that biological computation exhibits in some problems. These neuromorphic systems are organized as populations of excitatory and inhibitory spiking neurons with configurable synaptic connections (FACETS, 2005–2009; Navaridas et al., 2013; Benjamin et al., 2014; Merolla et al., 2014; Ning et al., 2015).

Synapses outnumber neurons by several orders of magnitude in biological neural networks (Binzegger et al., 2004). Reproducing these biological features in neuromorphic electronic circuits presents a scaling problem, as integrating thousands of dedicated synapse circuits per neuron can quickly become infeasible for systems that require a large number of neurons (Schemmel et al., 2007). This scaling problem has traditionally been solved by either treating synapses as simple linear elements and time-multiplexing spikes from many pre-synaptic

sources onto the same linear circuit (Benjamin et al., 2014), or by treating them as basic binary elements that can be set either $'on'$ or $'off'$ externally, without learning abilities (Merolla et al., 2014).

Real synapses, however, exhibit non-linear phenomena like spike timing dependent plasticity (STDP) that modulate the weight of an individual synapse based on the activity of the pre- and post-synaptic neurons (Bi and Poo, 1998). The modulation of synaptic weights through plasticity has been shown to greatly increase the range of computations that neural networks can perform (Abbott and Regehr, 2004). Capturing the plasticity properties of real synapses in analog neuromorphic hardware requires the use of distinct physical circuits/elements for each synapse. In conventional CMOS, this can lead to restrictions on scalability. Some potential solutions to the scalability issues in pure CMOS technology involve the use of very large integrated structures (e.g., up to a full wafer, Schemmel et al., 2012) or the adoption of deep submicron technologies (Noack et al., 2015). Scalability restrictions however can be greatly relaxed if one resorts to compact nano-scale circuit elements that can reproduce the plasticity properties of real synapses.

One potential candidate for these elements is the "memristor." Chua (1971) described the memristor as an element which *behaves somewhat like a non-linear resistor with memory.* Since HP first linked resistively switching devices with the concept of a memristor (Strukov et al., 2008), work on memristive devices has mostly focused on digital storage and logic functions (Linn et al., 2012; You et al., 2014), but there are also applications as analog/multi-level storage (Moreno et al., 2010; Shuai et al., 2013) and even memristive encryption (Lin and Wang, 2010; Du et al., 2014). In the neuromorphic community, memristors are seen as ideal devices for synapse implementations, as they combine three key functions in one device. Memristors can implement biologically realistic synaptic weight updates, i.e., learning (Jo et al., 2010), they can carry out long term multi-valued weight storage, and they can also communicate weighted pre-synaptic activity to the postsynaptic side (Saighi et al., 2015), significantly relaxing scalability restrictions (Indiveri et al., 2013).

Typically, plasticity in these memristive synapses is evoked by applying specific waveforms to the two terminals of the memristor, with the waveforms aligned to pre- respectively postsynaptic pulses (Jo et al., 2010). The correlation of the waveforms across the memristor in turn implements STDP-like plasticity (Mayr et al., 2012), with the form of the STDP curve defined by the applied wave shape (Serrano-Gotarredona et al., 2013). Both hardware and software models of plasticity based on the basic STDP mechanism are typically chosen primarily for their simplicity (Mayr and Partzsch, 2010). It has been argued however that more elaborate models of plasticity are required to reproduce the experimental evidence obtained from more complex synaptic plasticity experiments in real neural systems, and to implement algorithms that can learn to store and classify correlated patterns (Senn and Fusi, 2005; Sjöström et al., 2008; Lisman and Spruston, 2010).

In this work we present a neuromorphic implementation of one of these extended plasticity models that implements a spike-based Perceptron learning algorithm (Brader et al., 2007), which makes use of both analog CMOS circuits and TiO_{2-x} memristive devices. Compared to the more widely used STDP paradigm, the implementation of this learning algorithm on memristors does not employ the postsynaptic spike timing. Instead, it relies on the correlation of presynaptic spikes with signals derived from the postsynaptic neuron, such as its membrane potential and a measurement of its recent spiking activity. These requirements lead to a novel and quite different approach to the CMOS driver circuits which does not require the generation of temporally long waveforms on the pre- or postsynaptic sides.

In addition to spike timing, plasticity in biological synapses also depends on the firing rate of the post-synaptic neuron (Sjöström et al., 2001), a phenomenon that can not be captured by pair-wise STDP mechanisms (Pfister et al., 2006). The spike-based perceptron learning rule explicitly contains a term that reflects the recent firing rate of the neuron and is thus able to realize the rate-dependence of synaptic weight updates. The rule is also able to realize weight updates that depend on pre-post spike timing even though it does not explicitly depend on the post-synaptic spike times. Instead, it uses the membrane potential of the post-synaptic neuron as an indirect estimator of post-synaptic firing times. The rule is thus able to reasonably match the behavior of biological synapses while having a functional form that can be implemented efficiently on pure CMOS or on hybrid CMOS-memristor neuromorphic systems.

We introduce the spike-based Perceptron learning model in Section 2.1 and the TiO_{2-x} memristive devices employed in this implementation in Section 2.2. The adaptation of the learning model to memristors is described in Section 2.3. Considerations for crossbar operation of this paradigm are given in Section 2.4. Section 3.1 shows basic results characterizing operation of the memristors. Characterization of the learning CMOS driver circuits implemented in VLSI are detailed in Section 3.2. Finally, results from implementing the spike-based Perceptron learning with the CMOS driver circuits on the memristors are presented in Section 3.3.

2. Materials and Methods

2.1. The Plasticity Model

The spike-based Perceptron learning model of long-term plasticity has been introduced in Brader et al. (2007) based on earlier work in Fusi et al. (2000). The model represents a synapse with two stable weights, potentiated and depressed, whereby the transition between the two stable weights is done in an analog or graded manner. The synaptic weight $X(t)$ is influenced by a combination of pre- and post-synaptic activity, namely the pre-synaptic spike time t_{pre} and the value of the post-synaptic neuron membrane voltage $V_{mem}(t)$ and intra-cellular calcium concentration $C(t)$. A pre-synaptic spike arriving at t_{pre} reads the instantaneous post-synaptic values $V_{mem}(t_{pre})$ and $C(t_{pre})$. The change in $X(t)$ depends on these instantaneous values in the following way:

$$X \rightarrow X + a \quad if \quad \{V_{mem}(t_{pre}) > \theta_V \quad and \quad (1)$$

$$\theta_{up}^l < C(t_{pre}) < \theta_{up}^h\}$$

$$X \rightarrow X - b \quad if \quad \{V_{mem}(t_{pre}) \leq \theta_V \quad and \quad (2)$$

$$\theta_{down}^l < C(t_{pre}) < \theta_{down}^h\},$$

where a and b are jump sizes and θ_V is a voltage threshold. In other words, $X(t)$ is increased if $V_{mem}(t)$ is elevated (above θ_V) when the pre-synaptic spike arrives and decreased when $V_{mem}(t)$ is low at time t_{pre} provided that the calcium variable $C(t)$ is in the correct range. θ_{up}^l, θ_{up}^h, θ_{down}^l, and θ_{down}^h are thresholds on the calcium variable. The calcium variable $C(t)$ is an auxiliary variable that is a low-pass filtered version of the post-synaptic spikes (see Brader et al., 2007, for details). The variable $C(t)$ is incremented by J_C at each post-synaptic spike time t_i, where J_C reflects the magnitude of spike-triggered calcium influx into the cell. $C(t)$ decays with a time constant τ_C:

$$\frac{dC(t)}{dt} = -\frac{1}{\tau_C}C(t) + J_C \sum_i \delta(t - t_i) \quad (3)$$

The dependence of the weight updates on $C(t)$ allows the learning rule to enable/disable the weight updates based on the long-term average of post-synaptic activity. $X(t)$ continuously drifts toward one of two stable values based on whether it is above or below the threshold θ_X:

$$\frac{dX}{dt} = \alpha \qquad if \qquad X > \theta_X \quad (4)$$

$$\frac{dX}{dt} = -\beta \qquad if \qquad X \leq \theta_X \quad (5)$$

The weight $X(t)$ is bounded above and below by the two stable states X_{high} and X_{low} which are not shown in the equations to simplify the notation. **Figure 1** illustrates the relevant waveforms and parameters of the spike-based Perceptron learning rule.

The dynamics of the membrane potential variable, V_{mem}, which is used in Equations (1) and (2), depend on the neuron model used. The original neuron model used with the perceptron-learning rule is the simple constant leak integrate and fire neuron model (Brader et al., 2007). The neuron circuit we have in our neuromorphic chip, however, implements the more realistic adaptive exponential integrate and fire neuron model. This neuron circuit and the underlying model are described in detail in Indiveri et al. (2011) and Ning et al. (2015). The interaction between this adaptive exponential integrate and fire silicon neuron and the spike-based perceptron-learning rule is described in Indiveri et al. (2010).

Although the spike-based plasticity rule described above has been shown to reproduce, on average, the classical STDP phenomenology (Brader et al., 2007), it differs from the vast majority of spike-timing plasticity rules in that it does not explicitly depend on the precise timing of both pre- and post-synaptic neuron spikes. The compatibility with the classical STDP learning rule comes about through the rule's dependence on the post-synaptic neuron's membrane potential: a pre-synaptic spike that occurs when the post-synaptic membrane potential
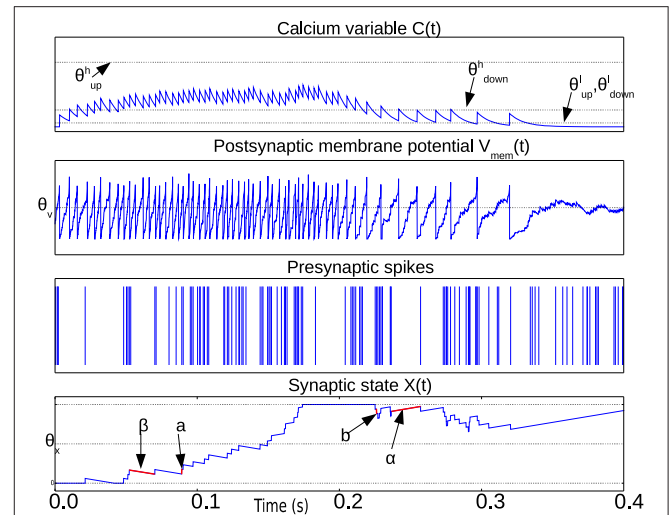
FIGURE 1 | Illustration waveforms of the spike-based perceptron learning rule showing key parameters from Equations (1–5). The Calcium variable plot shows the ranges defined by $\theta_{up}^l, \theta_{up}^h, \theta_{down}^l, \theta_{down}^h$ within which synaptic plasticity is active according to Equations (1) and (2). The post-synaptic neuron membrane potential plot shows the threshold θ_V. Incoming synapses can be depressed (potentiated) if $V_{mem}(t)$ is below (above) θ_V. The bottom plot showing the synaptic state $X(t)$ illustrates the jump and drift mechanism. On each pre-synaptic spike, the mutually exclusive conditions in Equations (1) and (2) are evaluated. If the condition in Equations (1) and (2) is fulfilled, the synaptic state jumps up (down) by a step a (b). The synaptic state is continuously drifting to a high or low state depending on whether it is above or below the threshold θ_X, respectively.

is high will potentiate the synapse and will likely produce a post-synaptic spike shortly after. Thus, the synapse tends to get potentiated in pre-before-post scenarios. The synapse also tends to get depressed in post-before-pre scenarios because the membrane potential is usually low for a few milliseconds after a post-synaptic spike is emitted, and a pre-synaptic spike arriving in this interval will depress the synapse.

The spike-based Perceptron plasticity rule also has access to post-synaptic neuron's rate information through the $C(t)$ signal. This allows it to reproduce effects beyond classical pair-wise STDP such as increased potentiation at high post-synaptic firing rates and increased depression at low post-synaptic firing rates (Sjöström et al., 2001). These effects arise in more complicated STDP models such as triplet STDP (Pfister et al., 2006; Mayr and Partzsch, 2010). The absence of explicit dependence on the post-synaptic neuron's firing times thus does not diminish the biological plausibility or the computational power of the spike-based Perceptron learning rule.

For the purpose of pure CMOS VLSI implementation (Chicca et al., 2014), this plasticity model is interesting because it can learn a graded response to an input pattern but on long time scales, the weight $X(t)$ drifts to one of two stable states and is thus easy to store long-term. In the hybrid CMOS-memristor architecture that we propose in this paper, however, the weight drift (Equations 4 and 5) is not implemented. The memristor conductance (weight) only changes on pre-synaptic spikes.

Weight drift or the bi-stable synaptic dynamics of the perceptron learning rule can be useful in consolidating the synaptic changes and making the synaptic weight more robust against spurious spikes (Brader et al., 2007). However, this comes at the cost of the sensitivity of the plasticity rule to the temporal spike patterns as multiple spike patterns might lead to the same binary synaptic weights. In the absence of weight drift as in the proposed hybrid CMOS-memristor architecture, the analog synaptic weights are able to maintain a synaptic trace that better reflects the identity of past spiking patterns (Maass and Markram, 2002).

2.2. Memristive Devices

The memristors that we use as synaptic elements are TiO_{2-x}-based memristors which were fabricated as follows: thermal oxidation was used to grow a 200 nm film of insulating SiO_2 on a 6″ Silicon wafer. Then, bottom electrodes (BEs) were patterned and obtained by conventional optical photolithography, electron beam evaporation and lift-off process. BEs consisted of evaporation of 5 nm adhesive Titanium (Ti) and 10 nm Platinum (Pt) layers. After that, a similar patterning process was used for the 25 nm TiO_{2-x} active layer that was deposited in a Leybold Helios Pro XL Sputterer to achieve high quality film. The film was sputtered from a Titanium metal target with 8 sccm flow of O_2, 35 sccm Ar, 2 kW at the cathode, and 15 sccm O_2, 2 kW at an additional plasma source. Then, again optical photolithography, electron beam evaporation and lift off process were used to pattern and deposit the 10 nm Pt top electrodes (TEs). **Figure 2** shows a cross-section and microphotograph of $Ti/Pt/TiO_{2-x}/Pt$ memristor prototype (device area: $60 \times 60\,\mu$m).
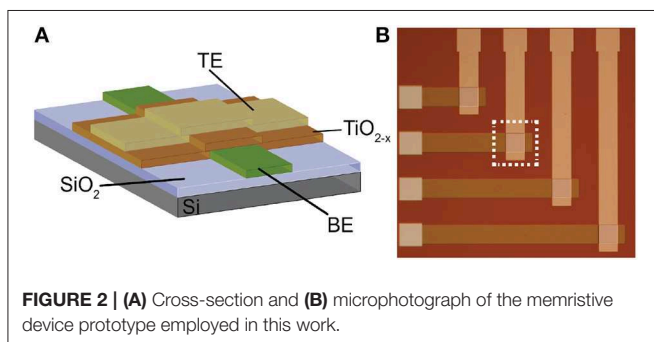
2.3. Circuits for Memristive Learning

The basic building block of the CMOS circuits is a neuron tile which is shown schematically in **Figure 3A**. The tile contains an analog subthreshold leaky integrate and fire neuron which is fully described in Qiao et al. (2015). The neuron integrates synaptic current (with an adjustable leak) on a capacitor. When the capacitor voltage crosses the firing threshold, the neuron generates a digital spike and the capacitor voltage is reset to ground. The plasticity circuit monitors the membrane potential, $Vmem$, and the spike output of the neuron and uses them to evaluate the conditions in Equations (1) and (2). The plasticity circuit internally generates the $C(t)$ signal by low-pass filtering the neuron spikes. The plasticity circuit then generates two digital signals: $'up'$ and $'dn'$ that determine whether incoming
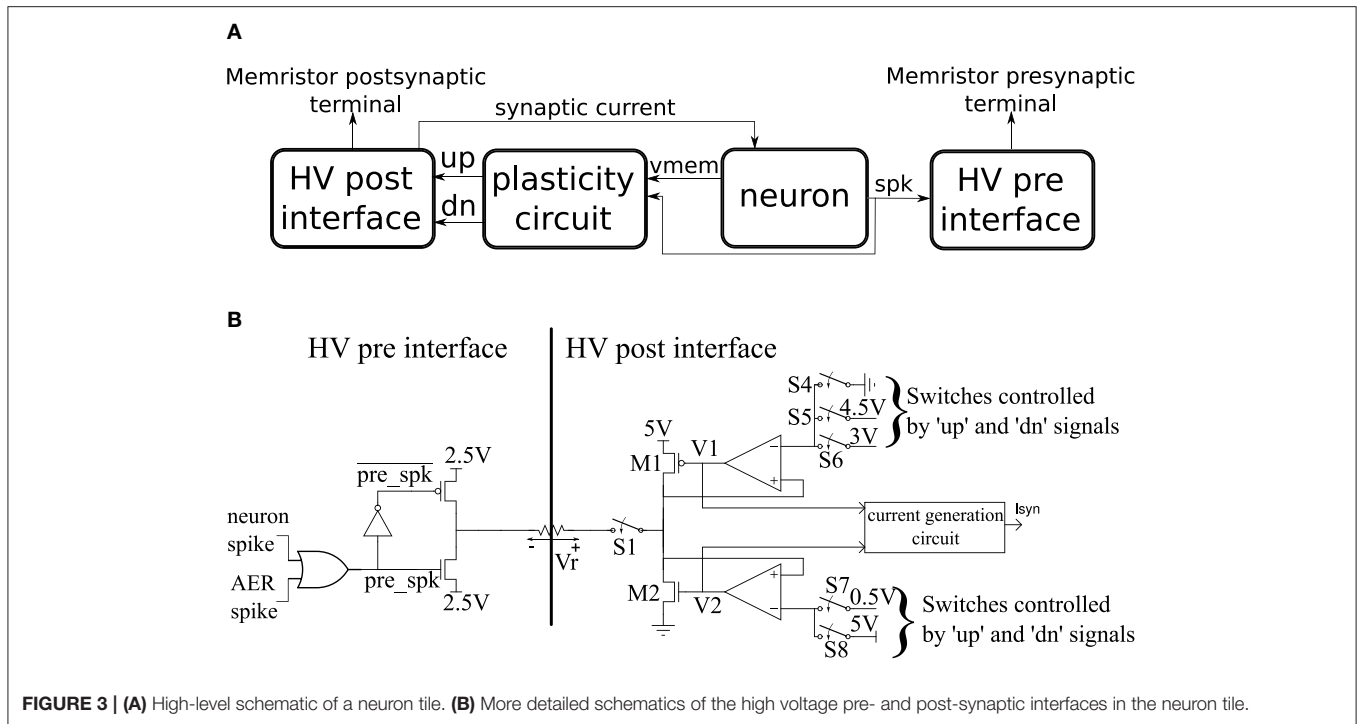
synapses/memristors should be potentiated, depressed, or left unchanged when a pre-synaptic spike occurs according to Equations (1) and (2). The plasticity circuit is described in more detail in Qiao et al. (2015).

A neuron tile has a pre-synaptic and a post-synaptic memristor terminal. These terminals are monitored and driven by the high voltage post- and pre- interfaces which run at a supply voltage of 5 V. All other circuits operate using a 1.8 V supply. The 5 V operation allows the memristor interface circuits to apply higher voltage pulses to the memristor terminals. The memristor conductance changes if pulses above a certain magnitude (the write threshold) are applied across it. The direction of the change depends on the polarity of the pulse. We designed the interface circuits so that they can interface to memristors having resistance values as low as 1 KOhm and deliver write pulses of either polarity with an amplitude of up to 2 V. The write voltage threshold for the memristor devices we use in this paper is much lower than 2 V. The height of the write pulses are programmable, however, so we can control their amplitudes up to 2 V. The width of the programming pulse is also configurable and can be as wide as 1 ms. The read pulse amplitude (which needs to be below the write threshold) is adjustable in the 0–2 V range and its width is also adjustable. The memristor is inserted between the pre-synaptic terminal of one tile and the post-synaptic terminal of another (or the same) tile. Spikes generated in the neuron circuit of the pre-synaptic tile will then cause a current proportional to the memristor conductance to be injected into the post-synaptic tile neuron. Moreover, based on the output of the plasticity circuit in the post-synaptic tile, a voltage pulse of the appropriate polarity is applied across the memristor terminals to increase/potentiate or decrease/depress its conductance when the pre-synaptic neuron tile generates a spike. In the rest of this section, we describe how this behavior is realized.

The pre- and post-synaptic memristor interfaces are shown in more detail in **Figure 3B** where they are linked by a memristive element. We retain the ability to disconnect the post-synaptic circuit from the memristor post-synaptic terminal using switch S1. The pre-synaptic memristor terminal is kept floating by default so no current can flow through the memristor and its value remains constant. The post-synaptic terminal is monitoring the current flowing through the memristor and injecting a proportional current into the neuron. By keeping the pre-synaptic terminal floating, no current flows through the memristor, and no current is injected into the post-synaptic neuron. When the pre-synaptic tile neuron spikes, or when the tile receives an AER event from off-chip, the pre-synaptic terminal is strongly clamped at 2.5 V for a short duration that is controlled by an analog bias. By clamping the pre-synaptic terminal to the middle of the supply voltage, we are able to apply pulses of either polarity with an amplitude of up to 2.5 V by setting the appropriate voltage on the post-synaptic terminal. If the post-synaptic terminal is clamped to V_{post}, then on pre-synaptic spike, a pulse of amplitude $V_{post} - 2.5$ is applied across the memristor. Assume switch S1 is closed. The post-synaptic terminal can be clamped to one of three possible values: 4.5, 0.5, or 3 V. These clamping voltages can be adjusted through analog biases. The clamping is done by the strong transistors M1 and

FIGURE 2 | (A) Cross-section and **(B)** microphotograph of the memristive device prototype employed in this work.

FIGURE 3 | (A) High-level schematic of a neuron tile. **(B)** More detailed schematics of the high voltage pre- and post-synaptic interfaces in the neuron tile.

M2 which are each part of a negative feedback loop that controls their gate potentials so as to maintain their drain potentials at one of the three voltage clamp values. A number of switches which are controlled by the $'up'$ and $'dn'$ signals from the plasticity circuit determine which clamping voltage is selected according to **Table 1**. For example, if switches S5 and S8 are closed and switches S4, S6, and S7 are open, the post-synaptic terminal is clamped by a PFET at 4.5 V. Switches S4–S8 are implemented as single transistors as each switch has to pass a bias voltages that is always either above 2.5 (PFET is used) or below 2.5 (NFET is used). Switch S1 is implemented as a transmission gate.

At a pre-synaptic spike which causes the pre-synaptic terminal to be clamped to 2.5 V, the memristor experiences a voltage pulse of either 2.0, −2.0, or 0.5 V depending on whether the post-synaptic terminal is at 4.5, 0.5, or 3V respectively. These three cases can either potentiate/increase the memristor conductance, depress/decrease it, or leave it unchanged respectively. It is the plasticity circuit, which through the $'up'$ and $'dn'$ signals controls switches S4–S8, which chooses between these three cases (**Table 1**).

The post-synaptic side indirectly senses the memristor conductance from the gate voltages V1 and V2. When the pre-synaptic side is floating, the two feedback loops push V1 and V2 to 5V and 0V, respectively. The current generation circuit will then generate very little current. At a pre-synaptic event, either V1 or V2 abruptly changes so that the actively clamping transistor has increased effective gate-source voltage so as to be able to source/sink the memristor current while maintaining the drain terminal at the clamp voltage. Larger memristor conductance translates to a larger change in V1 or V2 and based on this change, a proportional current $Isyn$ is generated and injected into the post-synaptic neuron. The current generation circuit

TABLE 1 | Effect of $'up'$ and $'dn'$ signals on the post-synaptic terminal potential which in turn determines the type of plasticity event induced on pre-synaptic spikes.

Plasticity signal	Vpost	Plasticity event	Open switches	Closed switches
$'up'$ = 0 and $'dn'$ = 0	3.0	No change	S4, S5, S7	S6, S8
$'up'$ = 1 and $'dn'$ = 0	4.5	Potentiate	S4, S6, S7	S5, S8
$'up'$ = 0 and $'dn'$ = 1	0.5	Depress	S8, S5, S6	S4, S7

Shown are the open and closed switches in each case. The switches are controlled by the $'up'$ and $'dn'$ signals.

approximately implements the equation:

$$Isyn = A * V2 - B * V1 \qquad (6)$$

Where A and B are constants adjusted through biases. This linear equation is, however, valid in a limited regime of V_1 and V_2. This regime can be adjusted through biases. Note that $Isyn$ is proportional to the absolute value of the memristor current, regardless of whether the current is sourced by transistor M1 or sunk by transistor M2. The $'up'$ and $'dn'$ signals can not both be high at the same time. For the three possible configurations of the $'up'$ and $'dn'$ signals in **Table 1**, M1 and M2 can not be supplying current at the same time. For the possible configuration of switches shown in **Table 1**, the feedback loops controlling V1 and V2 ensure that the gate-source voltage of M1 and M2 can not be simultaneously non-zero. This guarantees that the current in either M1 or M2 is the current flowing through the memristor.

This active clamp technique allows maximum voltage headroom for transistors M1 and M2 which allows them to clamp the post-synaptic terminal at voltages near the supply rails. It
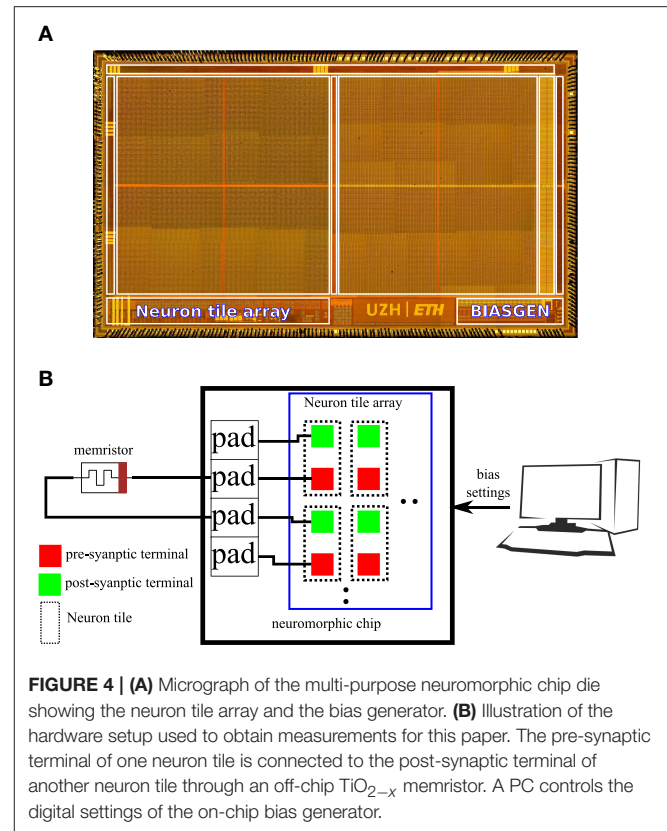
also enables precise control over the magnitude of the voltage pulses applied across the memristor. In Section 3.2, we present experimental results to illustrate the behavior of the circuit in **Figure 3A**.

The synaptic weight in the original spike-based Perceptron learning rule has only two stable states due to the weight drift (Equations 4 and 5) which pushes the weight to either a high or a low value. This mechanism is not present in our architecture; the synaptic weight (memristor conductance) is an analog quantity that only changes in response to pre-synaptic spikes and is stable otherwise. Realizing analog synaptic weights that are long-term stable is difficult in pure CMOS as analog weights that are encoded using charge on a capacitor are easily corruptible through leakage paths and capacitive coupling to nearby nodes. Therefore, in pure CMOS, a multi-stability mechanism is required to push the weights to well-defined and stable discrete states. Hybrid CMOS-memristor architectures like ours can realize naturally stable analog weights (memristor conductances) and thus do not require such a mechanism.

An 8 * 8 array of the neuron tile shown in **Figure 3** was fabricated on a standard 6 M 180 nm CMOS process as part of a larger multi-purpose neuromorphic chip shown in **Figure 4A**. The chip contains a bias generator based on the design in Delbruck and Lichtsteiner (2006). The bias generator has a low pin-count (5 pins) digital programming interface that can be used to set the values of the analog biases used in the neuron tile array. The other components of the multi-purpose neuromorphic chip are described in detail in Qiao et al. (2015) but they are not relevant for the current paper. Address event representation (AER) interfaces carry spikes to/from the neuron tile array. The pre- and post-synaptic terminals of the 64 neuron tiles were routed to the top-metal level to make it possible to directly deposit a cross-bar array of TiO_{2-x} memristors on top that connects a memristor between each pre-synaptic terminal and each post-synaptic terminal. This post-processing step was not carried out. In the chip, the pre- and post-synaptic memristor terminals of two neuron tiles were directly connected to pads. An off-chip memristor was then connected between the pre-synaptic terminal of one of these neuron tiles and the post-synaptic terminal of the other tile as shown in **Figure 4B**. This setup was used to obtain the measurements presented in the rest of this paper.

2.4. Crossbar Operation

In a crossbar configuration, N neuron tiles are interconnected by an array of N^2 memristors where there is a memristor connecting each pre-synaptic terminal to each post-synaptic terminal. To achieve high synaptic/memristor integration densities, it is important to avoid memristor specific CMOS circuits and only access the memristors through the N pre-synaptic terminals and N postsynaptic terminals which form the row lines and column lines of the $N * N$ memristor array. Consider the simple case of $N = 2$ neuron tiles connected using $N^2 = 4$ memristors. If the post-interfaces in tile 1 and 2 are clamping the post-synaptic terminals to different voltages (which would be the case if one of them is in the ′up′ state and the other is in the ′dn′ state) then current would flow between the post-synaptic terminals through

FIGURE 4 | (A) Micrograph of the multi-purpose neuromorphic chip die showing the neuron tile array and the bias generator. **(B)** Illustration of the hardware setup used to obtain measurements for this paper. The pre-synaptic terminal of one neuron tile is connected to the post-synaptic terminal of another neuron tile through an off-chip TiO_{2-x} memristor. A PC controls the digital settings of the on-chip bias generator.

two memristors connected in series. This would lead to changes in the conductances of these memristors in the absence of pre-synaptic spikes and to synaptic current being mistakenly injected into the neurons. Crossbar operation is thus only possible if plasticity is switched off through the analog biases so that all post-synaptic terminals are clamped at the same potential.

One benefit of using a crossbar array to implement synaptic matrices is that a post-synaptic neuron only needs to know the aggregate input it receives from all synapses rather than the individual contributions. This considerably relaxes the design of the driver circuits as these circuits do not need to isolate the contribution of single devices. Moreover, it has been shown that small selector-less arrays can perform quite well even as analog memory, where good isolation of the contribution of each individual element is required, provided certain assumptions about the switching characteristics of the memristors (e.g., concerning maximum and minimum resistive states) hold (Serb et al., 2015). Thus, even though the current implementation does not take the additional complications of crossbar configurations into account, there is evidence that extension of our work to first small, selector-less arrays and then potentially to larger selector-based arrays is possible.

3. Results

3.1. Initial Memristor Programming Results

Before the memristors could be used as artificial synapses they were electrically prepared for operation. The preparation

procedure consisted of an electroforming step, a stabilization period and a characterization stage. During all stages the devices are biased by a series of square-wave pulses of fixed duration (100 μ s) and variable amplitude.

Initially, the measured resistive state (RS) of all our devices was above the 10 MΩ mark. During electroforming devices were subjected to voltage pulse ramps beginning at 1 V and increasing in steps of 0.5 V until the RS dropped to below 500 kΩ or the maximum limit of 8 V was reached. Typically, electroforming was achieved after applying a 6 V pulse. During electroforming, voltage was applied to our devices through a 100 kΩ series resistor as a measure to protect them against unduly high power dissipation and consequent damage.

In the stabilization period devices are subjected to pulse trains whose amplitudes and polarities are determined by trial and error. During this phase the devices are forced to oscillate between more resistive and more conductive states. This is achieved through use of a bipolar stimulation protocol, that is pulses of opposite polarities drive the RS in opposite directions. During the stabilization period no stable limits for the operational RS ceiling and floor can be reliably determined, nor can appropriate voltages be found at which the memristor will reliably switch between floor and ceiling.

The characterization stage follows seamlessly from the stabilization period as the device settles to an operational RS range. In this phase voltage pulse amplitudes are trimmed until a set of amplitudes for normal operation biasing is selected. **Figure 5** shows a typical characterization stage series of read-outs obtained from a well-behaved device. Notably different voltage values are tested for both SET-type (toward lower RS) and RESET-type pulse polarities within a relatively narrow range ($\approx$ 200 mV). Typically devices can operate comfortably within such narrow ranges although their operational range and the number of pulses it takes to transition between floor and ceiling (and vice versa) will be affected by the exact choice of pulse voltage. See **Figure 5** for an example.

Once pulsing voltages have been determined, the memristor may be connected to the appropriately configured neuromorphic circuitry, ready for bipolar-mode operation.

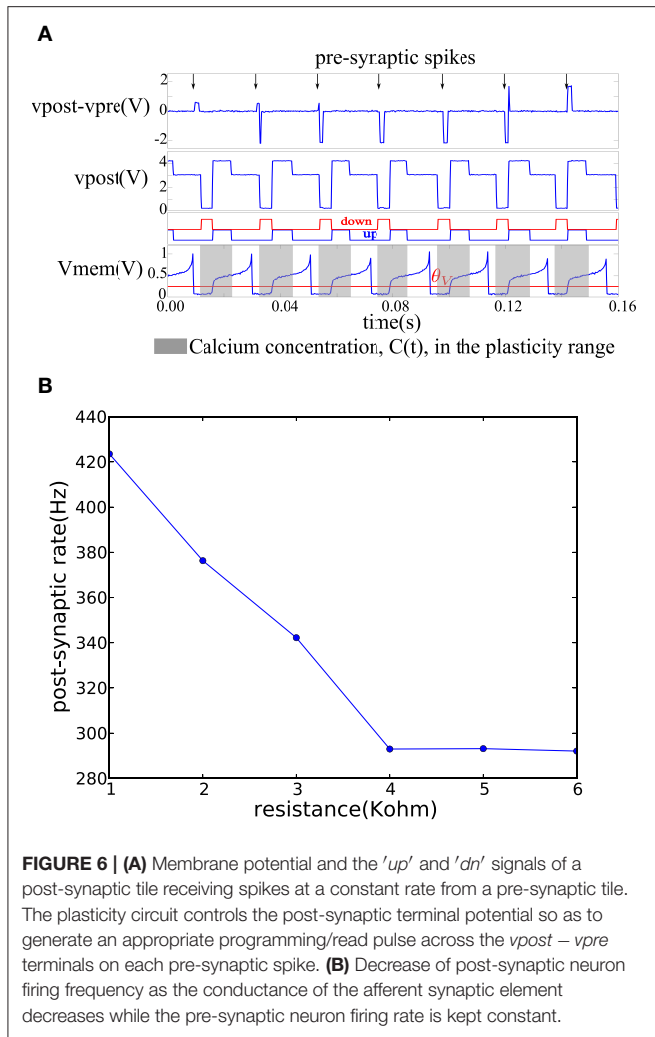## 3.2. Characterization of CMOS Plasticity Circuits

On each pre-synaptic spike, the pre- and post-synaptic interface circuits in the neuron tile shown in **Figure 3** apply a voltage pulse to update the memristor value according to the spike-based perceptron learning rule. This behavior is illustrated in **Figure 6A** where a fixed resistor was inserted between the pre-synaptic terminal of one tile and the post-synaptic terminal of another (as in **Figure 4B** but using a resistor instead of a memristor). Constant current is injected into the neurons to maintain a constant firing rate. The calcium signal, $C(t)$, jumps up after each spike and enters the plasticity range, then it decays back out of the plasticity range. The bottom two plots in **Figure 6A** show the $'up'$ and $'dn'$ signals and the neuron membrane potential in the post-synaptic tile. The $'up'$ and $'dn'$ signals are generated by the plasticity circuit in the post-synaptic tile (see **Figure 3A**). This plasticity circuit calculates the calcium variable, $C(t)$, from the post-synaptic neuron spikes according to Equation (3). It
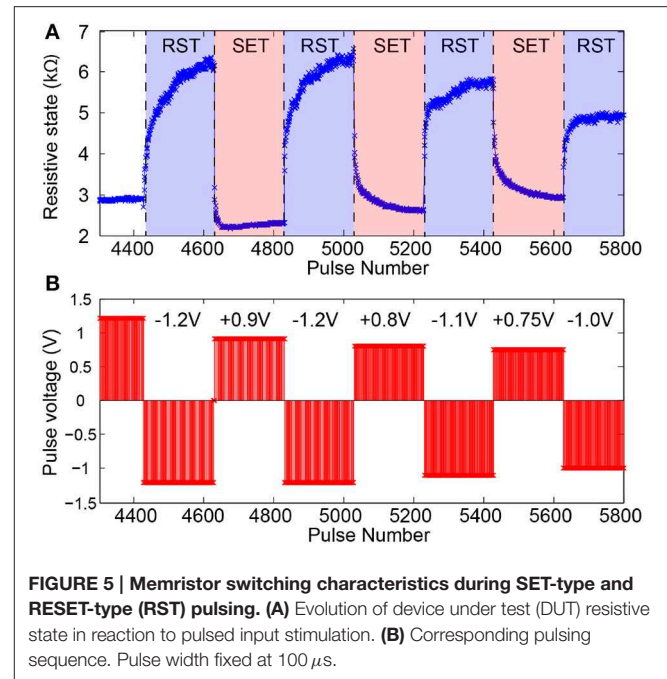
FIGURE 5 | Memristor switching characteristics during SET-type and RESET-type (RST) pulsing. (A) Evolution of device under test (DUT) resistive state in reaction to pulsed input stimulation. **(B)** Corresponding pulsing sequence. Pulse width fixed at 100 μs.

evaluates the conditions in Equations (1) and (2) to decide whether to potentiate, depress, or leave unchanged incoming synapses when the pre-synaptic neuron spikes. This decision is communicated to the post-synaptic interface circuit which clamps the post-synaptic terminal voltage $vpost$ at 4.1V (when the $'up'$ signal is high), 0.1 V (when the $'dn'$ signal is high), or 3 V (when both the $'dn'$ and $'up'$ are low) as shown in **Figure 6A**. The pre-synaptic terminal is floating by default and is clamped at 2.5 V for a short duration on each pre-synaptic spike. For each pre-synaptic spike, this causes $vpost - vpre$ to be approximately 2 V when the $'up'$ signal is high which would increase the memristor conductance (potentiation), –2 V when the $'dn'$ signal is high which would decrease the memristor conductance (depression), and 0.5 V otherwise as shown in **Figure 6A** which would leave the memristor conductance unchanged and simply read out its value. In **Figure 6A**, at the first pre-synaptic spike, $C(t)$ is outside the plasticity range and a small read pulse is applied. The subsequent pre-synaptic spikes arrive first in the depression, then in the potentiation intervals of the post-synaptic tile and large amplitude pulses with the appropriate polarity are applied.

Figure 6B shows how the firing frequency of the post-synaptic neuron varies as a function of the value of the resistor connecting it to the pre-synaptic tile. The pre-synaptic tile generates spikes at a constant frequency. The firing frequency of the post-synaptic neuron steadily decreases with the decreasing conductance of the resistor. The bias conditions were chosen to obtain a linear region in the 1–4 KOhm resistance range beyond which the post-synaptic neuron firing frequency saturates at a lower bound. The neuron is biased to have a spontaneous baseline firing rate which is about 290 Hz. The transfer function from the synaptic resistance to the synaptic current injected into the neuron is linear in the 1–4 KOhm region in **Figure 6B** but beyond that, it is highly non-linear causing a small increase in synaptic resistance

FIGURE 6 | (A) Membrane potential and the $'up'$ and $'dn'$ signals of a post-synaptic tile receiving spikes at a constant rate from a pre-synaptic tile. The plasticity circuit controls the post-synaptic terminal potential so as to generate an appropriate programming/read pulse across the $vpost - vpre$ terminals on each pre-synaptic spike. **(B)** Decrease of post-synaptic neuron firing frequency as the conductance of the afferent synaptic element decreases while the pre-synaptic neuron firing rate is kept constant.

FIGURE 7 | (A) The post-synaptic membrane potential in response to a pre-synaptic pulse train of 47 Hz. Vertical lines indicate the times of two pre-synaptic spikes. **(B)** Frequency of the post-synaptic neuron after a number of plasticity episodes. Plasticity episodes were applied in succession by adjusting the analog biases to put the post-synaptic neuron tile in the potentiation and depression regimes.

to lead to a greatly reduced synaptic current which becomes negligible compared to the constant injection current used to maintain the baseline firing rate. The neuron then saturates at the baseline firing rate.

3.3. Memristive Plasticity Experiments

In the fabricated chip, the pre- and post-synaptic terminals of two neuron tiles were available on the chip pads. The pre-synaptic terminal of one tile was connected to the post-synaptic terminal of the other tile through a TiO_{2-x} memristor as shown in **Figure 4B**. The pre-synaptic neuron was biased to fire at 47 Hz. The spike-based perceptron learning circuit was then successively cycled between the potentiation and depression regimes (using the analog biases) and the resulting post-synaptic firing rate was observed. The post-synaptic firing rate was taken as an indication of the synaptic weight or the conductance of the memristive element. The system showed correct operation with the postsynaptic firing rate increasing after a potentiation episode and decreasing after a depression episode as shown in **Figure 7**.

A potentiation episode involves setting the biases of the plasticity circuit in **Figure 3A** so that its $'up'$ output signal is constantly high which will cause the memristor post-synaptic
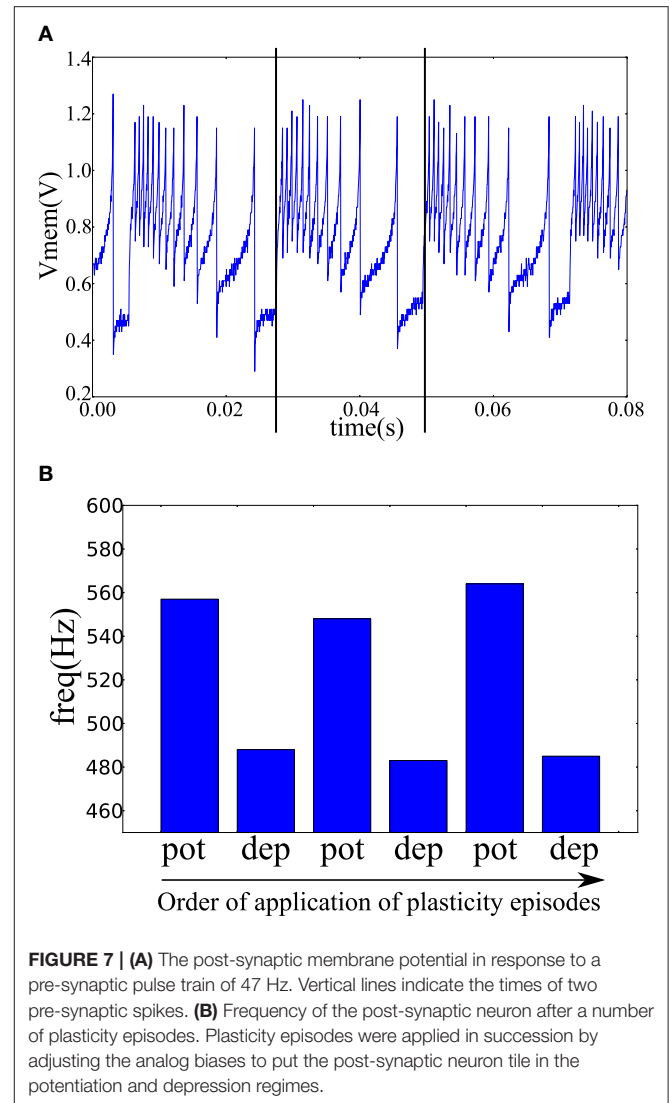
terminal to be clamped at approximately 4.1 V as shown in **Figure 6A**. On each pre-synaptic spike, a pulse of approximately 2 V is thus applied across the memristor terminals which will act to increase its conductance. Similarly, in a depression episode, the $'dn'$ output signal is constantly high which will cause a pulse of approximately –2 V to be applied across the memristor terminals on each pre-synaptic spike which will act to decrease its conductance. After each plasticity episode, plasticity was disabled and the post-synaptic neuron firing rate measured, then the next plasticity episode is applied.

4. Discussion

4.1. Memristive Device Characteristics

Figure 5 shows the typical operation of a "well-behaved" memristor in response to trains of input voltage pulses. A number of key features are noteworthy:

- *Bipolar operation:* Pulses of opposite polarity precipitate resistive state changes in opposite directions. In the case of our

devices, a positive voltage applied to the top electrode (bottom electrode grounded) causes potentiation.

- *Bidirectionally gradual switching:* Transitions between resistive state floor and ceiling occur over many pulses, not just one. This allows the device to work as a multi-level weight artificial synapse (as opposed to binary).
- *Bidirectionally saturating switching:* When a device is bombarded by trains of identical voltage pulses it approaches its operational resistive state floor and ceiling in progressively smaller steps. This implies that the middle of the resistive state range is expected to be most often unoccupied *in operando*, as it is traversed quickly in either direction under pulsing. The resistive state will be therefore multi-level in nature, but most of the time distinctly high or low.
- *Biasing parameter variation tolerance:* The device can remain functional under a relatively wide range of bias voltages. We obtain good switching behavior for voltage pulses in the 0.75–1.2 V range. The device can safely operate with voltage pulses of up to 2 V. This bodes well for operation in tandem with practical electronic systems and for resistive switching behavior tuning.

These features allow the memristive devices to exhibit the correct behavior when coupled to the neuromorphic circuits described in Section 2.3, both as binary and as multi-level synapses. Only binary synaptic operation was investigated in the plasticity experiments.

4.2. The Spike-based Perceptron Learning Rule in CMOS-memristor Architectures

The spike-based Perceptron plasticity rule has been implemented in CMOS neuromorphic systems using various types of circuits such as subthreshold circuits (Mitra et al., 2009) and switched capacitor circuits (Noack et al., 2015). In this paper, we have presented a physical implementation of the first hybrid CMOS-memristor architecture that implements a spike-based Perceptron learning plasticity rule. The physical CMOS-memristor system we presented is a standalone system in which the custom CMOS chip connects directly to the memristive devices. The CMOS chip implements the neuron elements together with dedicated per-neuron circuits that can program (potentiate or depress) the memristive synaptic elements as well as sense their conductances/weights to generate proportional Excitatory Post-Synaptic Currents (EPSCs) in the post-synaptic neuron in response to pre-synaptic spikes. We have presented direct measurements that illustrate the behavior of this physical CMOS-memristor system. This is the first standalone neuromorphic system that combines custom neuron circuits with memristor programming and sensing circuits acting on physical memristive devices.

Many highly accurate and biologically grounded, i.e., non-empirical, synaptic plasticity rules make use of several auxiliary variables beyond spike times in the pre- and post-synaptic neurons to control synaptic weight updates (Pfister et al., 2006; Brader et al., 2007; Clopath and Gerstner, 2010; Mayr and Partzsch, 2010; Graupner and Brunel, 2012). These auxiliary variables may include low-pass filtered versions of the membrane potential (Clopath and Gerstner, 2010) or a low-pass filtered version of the neuron's spike train (Brader et al., 2007). Interestingly, the time difference between pre- and post-synaptic spikes does not figure explicitly in these models. This presents a problem for current neuromorphic memristive architectures that mainly depend on this time difference (through the overlap between pre- and post-synaptic spike-triggered waveforms) to induce weight updates. These architectures will not be able to handle weight updates that are triggered on single pre- or post-synaptic spikes.

The architecture we presented triggers weight updates on single pre-synaptic spikes. This has a significant advantage: at the time of a pre-synaptic spike, the neuromorphic synapse can be immediately potentiated or depressed based on the current state of the post-synaptic neuron; the neuromorphic system does not have to wait for a post-synaptic spike to know the outcome of the plasticity event. Implementations of classical pair-wise STDP rules using memristors typically trigger long waveforms on the pre- and post-synaptic sides of the memristor in response to pre- and post-synaptic spikes respectively. When these waveforms overlap, the potential difference across the memristor exceeds a threshold and changes in memristor conductance occur. The duration of these waveforms dictate the STDP window. The overlapping waveforms paradigm is problematic in the high spike rate regime as multiple spikes can occur within the STDP window, thereby corrupting the synaptic weight update. By contrast this problem is completely avoided in the case of the spike-based Perceptron learning rule.

In the original learning rule (Brader et al., 2007) the weights were bistable, i.e., they gradually drifted to one of two stable states. This had the effect of consolidating synaptic changes and making it more difficult for a synaptic pattern to be corrupted by spurious spikes. Our architecture does not implement such continuous (non event-driven) weight drift. This indicates that synaptic rule features that simplify pure CMOS implementations like bistable weights do not necessarily translate to simpler CMOS-memristor implementations.

4.3. Outlook

The architecture we describe represents a first step toward hybrid CMOS-memristor implementations of more elaborate plasticity rules that go beyond standard STDP. Further developments will have to address the problem of plastic crossbar operation as well as mechanisms that allow continuous or non event-driven weight updates.

Acknowledgment

References

Abbott, L. F., and Regehr, W. G. (2004). Synaptic computation. *Nature* 431, 796–803. doi: 10.1038/nature03010

Benjamin, B. V., Gao, P., McQuinn, E., Choudhary, S., Chandrasekaran A. R., Bussat, J., et al. (2014). Neurogrid: a mixed-analog-digital multichip system for large-scale neural simulations. *Proc. IEEE* 102, 699–716. doi: 10.1109/jproc.2014.2313565

Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.

Binzegger, T., Douglas, R. J., and Martin, K. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* 24, 8441–8453. doi: 10.1523/JNEUROSCI.1400-04.2004

Brader, J. M., Senn, W., and Fusi, S. (2007). Learning real world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* 19, 2881–2912. doi: 10.1162/neco.2007.19.11.2881

Chicca, E., Stefanini, F., Bartolozzi, C., and Indiveri, G. (2014). Neuromorphic electronic circuits for building autonomous cognitive systems. *Proc. IEEE* 102, 1367–1388. doi: 10.1109/JPROC.2014.2313954

Chua, L. (1971). Memristor-the missing circuit element. *Circ. Theory IEEE Trans.* 18, 507–519. doi: 10.1109/TCT.1971.1083337

Clopath, C., and Gerstner, W. (2010). Voltage and spike timing interact in stdp – a unified model. *Front. Synaptic Neurosci.* 2:25. doi: 10.3389/fnsyn.2010.00025

Delbruck, T., and Lichtsteiner, P. (2006). "Fully programmable bias current generator with 24 bit resolution per bias," in *Circuits and Systems, 2006. ISCAS 2006. Proceedings 2006. IEEE International Symposium* (island of Kos: IEEE), 4.

Du, N., Manjunath, N., Shuai, Y., Bürger, D., Skorupa, I., Schüffny, R., et al. (2014). Novel implementation of memristive systems for data encryption and obfuscation. *J. Appl. Phys.* 115, 124501. doi: 10.1063/1.4869262

FACETS. (2005–2009). *Fast Analog omputing with Emergent Transient States in Neural Architectures (FACETS).* FP6-2005-015879 EU Grant. Heidelberg.

Fusi, S., Annunziato, M., Badoni, D., Salamon, A., and Amit, D. (2000). Spike–driven synaptic plasticity: theory, simulation, VLSI implementation. *Neural Comput.* 12, 2227–2258. doi: 10.1162/089976600300014917

Graupner, M., and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proc. Natl. Acad. Sci. U.S.A.* 109, 3991–3996. doi: 10.1073/pnas.1109359109

Indiveri, G., Stefanini, F., and Chicca, E. (2010). "Spike-based learning with a generalized integrate and fire silicon neuron," in *International Symposium on Circuits and Systems, (ISCAS), 2010* (Paris: IEEE), 1951–1954. Available online at: http://ncs.ethz.ch/pubs/pdf/Indiveri_etal10.pdf

Indiveri, G., Linares-Barranco, B., Hamilton, T. J., van Schaik, A., Etienne-Cummings, R., Delbruck, T., et al. (2011). Neuromorphic silicon neuron circuits. *Front. Neurosci.* 5:73. doi: 10.3389/fnins.2011.00073

Indiveri, G., Linares-Barranco, B., Legenstein, R., Deligeorgis, G., and Prodromakis, T. (2013). Integration of nanoscale memristor synapses in neuromorphic computing architectures. *Nanotechnology* 24:384010. doi: 10.1088/0957-4484/24/38/384010

Jo, S. H., Chang, T., Ebong, I., Bhadviya, B. B., Mazumder, P., and Lu, W. (2010). Nanoscale memristor device as synapse in neuromorphic systems. *Nano Lett.* 10, 1297–1301. doi: 10.1021/nl904092h

Lin, Z., and Wang, H. (2010). Efficient image encryption using a chaos-based pwl memristor. *IETE Tech. Rev.* 27, 318–325. doi: 10.4103/0256-4602.64605

Linn, E., Rosezin, R., Tappertzhofen, S., Böttger, U., and Waser, R. (2012). Beyond von neumann - logic operations in passive crossbar arrays alongside memory perations. *Nanotechnology* 23:305205. doi: 10.1088/0957-4484/23/30/305205

Lisman, J., and Spruston, N. (2010). Questions about stdp as a general model of synaptic plasticity. *Front. Synaptic Neurosci.* 2:140. doi: 10.3389/fnsyn.2010.00140

Maass, W., and Markram, H. (2002). Synapses as dynamic memory buffers. *Neural Netw.* 15, 155–161. doi: 10.1016/S0893-6080(01)00144-7

Mayr, C. G., and Partzsch, J. (2010). Rate and pulse based plasticity governed by local synaptic state variables. *Front. Synaptic Neurosci.* 2:28. doi: 10.3389/fnsyn.2010.00033

Mayr, C., Stärke, P., Partzsch, J., Cederstroem, L., Schüffny, R., Shuai, Y., et al. (2012). "Waveform driven plasticity in BiFeO3 memristive devices: model and implementation," in *Advances in Neural Information Processing Systems 25* (Lake Tahoe, CA), 1700–1708.

Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642

Mitra, S., Fusi, S., and Indiveri, G. (2009). Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. *Biomed. Circ. Syst. IEEE Trans.* 3, 32–42. doi: 10.1109/TBCAS.2008.2005781

Moreno, C., Munuera, C., Valencia, S., Kronast, F., Obradors, X., and Ocal, C. (2010). Reversible resistive switching and multilevel recording in La0.7Sr0.3MnO3 thin films for low cost nonvolatile memories. *Nano Lett.* 10, 3828–3835. doi: 10.1021/nl1008162

Navaridas, J., Furber, S., Garside, J., Jin, X., Khan, M., Lester, D., et al. (2013). Spinnaker: fault tolerance in a power-and area-constrained large-scale neuromimetic architecture. *Parallel Comput.* 39, 693–708. doi: 10.1016/j.parco.2013.09.001

Ning, Q., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., et al. (2015). A re-configurable on-line learning spiking neuromorphic processor. *Front. Neurosci.* 9:141. doi: 10.3389/fnins.2015.00141

Noack, M., Partzsch, J., Mayr, C. G., Hänzsche, S., Scholze, S., Höppner, S., et al. (2015). Switched-capacitor realization of presynaptic short-term-plasticity and stop-learning synapses in 28 nm CMOS. *Front. Neurosci.* 9:10. doi: 10.3389/fnins.2015.00010

Pfister, J. P., Toyoizumi, T., Barber, D., and Gerstner, W. (2006). Optimal spike-timing dependent plasticity for precise action potential firing in supervised learning. *Neural Comput.* 18, 1309–1339. doi: 10.1162/neco.2006.18.6.1318

Qiao, N., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., et al. (2015). A re-configurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128k synapses. *Front. Neurosci.* 9:141. doi: 10.3389/fnins.2015.00141

Saighi, S., Mayr, C., Linares-Barranco, B., Serrano-Gotarredona, T., Schmidt, H., Lecerf, G., et al. (2015). Plasticity in memristive devices. *Front. Neurosci.* 9:51. doi: 10.3389/fnins.2015.00051

Schemmel, J., Brüderle, D., Meier, K., and Ostendorf, B. (2007). "Modeling synaptic plasticity within networks of highly accelerated I&F neurons," in *International Symposium on Circuits and Systems, (ISCAS), 2007* (New Orleans, LA: IEEE), 3367–3370.

Schemmel, J., Grübl, A., Hartmann, S., Kononov, A., Mayr, C., Meier, K., et al. (2012). "Live demonstration: a scaled-down version of the BrainScaleS wafer-scale neuromorphic system," in *IEEE International Symposium on Circuits and Systems ISCAS 2012* (Seoul), 702.

Senn, W., and Fusi, S. (2005). Learning only when necessary: better memories of correlated patterns in networks with bounded synapses. *Neural Comput.* 17, 2106–2138. doi: 10.1162/0899766054615644

Serb, A., Redman-White, W., Papavassiliou, C., Berdan, R., and Prodromakis, T. (2015). "Limitations and precision requirements for read-out of passive, linear, selectorless rram arrays," in *Circuits and Systems (ISCAS), 2015 IEEE International Symposium* (Lisbon: IEEE), 189–192.

Serrano-Gotarredona, T., Masquelier, T., Prodromakis, T., Indiveri, G., and Linares-Barranco, B. (2013). STDP and STDP variations with memristors for spiking neuromorphic learning systems. *Front. Neurosci.* 7:2. doi: 10.3389/fnins.2013.00002

Shuai, Y., Ou, X., Luo, W., Du, N., Wu, C., Zhang, W., et al. (2013). Nonvolatile multilevel resistive switching in Ar+ irradiated BiFeO3 thin films. *IEEE Electron Device Lett.* 34, 54–56. doi: 10.1109/LED.2012.2227666

Sjöström, P., Turrigiano, G., and Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6

Sjöström, P., Rancz, E., Roth, A., and Häusser, M. (2008). Dendritic excitability and synaptic plasticity. *Physiol. Rev.* 88, 769–840. doi: 10.1152/physrev.00016.2007

Strukov, D. B., Snider, G. S., Stewart, D. R., and Williams, R. S. (2008). The missing memristor found. *Nature* 453, 80–83. doi: 10.1038/nature06932

You, T., Shuai, Y., Luo, W., Du, N., Bürger, D., Skorupa, I., et al. (2014). Exploiting memristive BiFeO3 bilayer structures for compact sequential logics. *Adv. Funct. Mater.* 24, 3357–3365. doi: 10.1002/adfm.201303365