Low-Energy Deep Belief Networks Using Intrinsic Sigmoidal Spintronic-based Probabilistic Neurons

Ramtin Zand¹, Kerem Yunus Camsari², Steven D. Pyle¹, Ibrahim Ahmed³,
Chris H. Kim³, and Ronald F. DeMara¹

¹Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL, 32816 USA

²Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, 47906 USA

³Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA

ABSTRACT

A low-energy hardware implementation of deep belief network (DBN) architecture is developed using near-zero energy barrier probabilistic spin logic devices (p-bits), which are modeled to realize an intrinsic sigmoidal activation function. A CMOS/spin based weighted array structure is designed to implement a restricted Boltzmann machine (RBM). Device-level simulations based on precise physics relations are used to validate the sigmoidal relation between the output probability of a p-bit and its input currents. Characteristics of the resistive networks and p-bits are modeled in SPICE to perform a circuit-level simulation investigating the performance, area, and power consumption tradeoffs of the weighted array. In the application-level simulation, a DBN is implemented in MATLAB for digit recognition using the extracted device and circuit behavioral models. The MNIST data set is used to assess the accuracy of the DBN using 5,000 training images for five distinct network topologies. The results indicate that a baseline error rate of 36.8% for a 784×10 DBN trained by 100 samples can be reduced to only 3.7% using a 784×800×800×10 DBN trained by 5,000 input samples. Finally, Power dissipation and accuracy tradeoffs for probabilistic computing mechanisms using resistive devices are identified.

ACM Reference Format:

Ramtin Zand¹, Kerem Yunus Camsari², Steven D. Pyle¹, Ibrahim Ahmed³, Chris H. Kim³, and Ronald F. DeMara¹. 2018. Low-Energy Deep Belief Networks Using Intrinsic Sigmoidal Spintronic-based Probabilistic Neurons.

In *GLSVLSI '18: 2018 Great Lakes Symposium on VLSI, May 23–25, 2018, Chicago, IL, USA*. ACM, New York, NY, USA, Article 4, 6 pages.

1 INTRODUCTION

The interrelated fields of machine learning (ML), and artificial neural networks (ANN) have grown significantly in previous decades due to the availability of powerful computing systems to train and simulate large scale ANNs within reasonable time-scales, as well as the abundance of data available to train such networks in recent years. The resulting research has realized a bevy of ANN architectures that have performed incredible feats including a wide range of classification problems, and various recognition tasks.

Most ML techniques in-use today rely on supervised learning, where the systems are trained on patterns with a known desired output, or label. However, intelligent biological systems exhibit unsupervised learning whereby statistically correlated input modalities are associated within an internal model used for probabilistic inference and decision making [5]. One interesting class of unsupervised learning approaches that has been extensively researched is the Restricted Boltzmann machine (RBM) [11]. RBMs can be hierarchically organized to realize deep belief networks (DBNs) that have demonstrated unsupervised learning abilities, such as natural language understanding [17]. Most RBM and DBN research has focused on software implementations, which provides flexibility, but requires significant execution time and energy due to large matrix multiplications that are relatively inefficient when implemented on standard Von-Neumann architectures due to the memory-processor bandwidth bottleneck when compared to hardware-based in-memory computing approaches [16]. Thus, research into hardware-based RBM designs has recently sought to alleviate these constraints.

Previous hardware-based RBM implementations have aimed to overcome software limitations by utilizing FPGAs [12, 15] and stochastic CMOS [2]. In recent years, emerging technologies such as resistive RAM (RRAM) [4, 21] and phase change memory (PCM) [9] are proposed to be leveraged within the DBN architecture as weighted connections interconnecting building blocks in RBMs. While most of the previous hybrid Memristor/CMOS designs focus on improving the synapse behaviors, the work presented herein overcomes many of the preceding challenges by utilizing a novel spintronic p-bit device that leverages intrinsic thermal noise within low energy barrier nanomagnets to provide a natural building block for RBMs within a compact and low-energy package. The contribution of this paper is go to beyond using low-energy barrier magnetic tunnel junctions (MTJs), as has been previously introduced for a neuron in spiking neuromorphic systems [19, 20]. To the best of our knowledge this paper is the first effort to use MTJs with near-zero energy barriers as neurons within an RBM implementation. Additionally, various parameters of a hybrid CMOS/spin weight array structure are investigated for metrics of power dissipation, and error rate using the MNIST digit recognition benchmarks.

2 FUNDAMENTALS OF RBM

Boltzmann Machines (BM) are a class of recurrent stochastic ANNs with binary nodes whereby each possible state of the network, v , has an energy determined by the undirected connection weights between nodes and the node bias as described by (1), where s_i^v is the state of node i in v , b_i is the bias, or intrinsic excitability of node

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

GLSVLSI '18, May 23–25, 2018, Chicago, IL, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5724-1/18/05...\$15.00

<https://doi.org/10.1145/3194554.3194558>

i , and w_{ij} is the connection weight between nodes i and j [1].

$$E(v) = - \sum_i s_i^v b_i - \sum_{i < j} s_i^v s_j^v w_{ij} \quad (1)$$

$$P(s_i = 1) = \sigma(b_i + \sum_j w_{ij} s_j) \quad (2)$$

Each node in a BM has a probability to be in state 1 according to (2), where σ is the logistic sigmoid function. BMs, when given enough time, will reach a Boltzmann distribution where the probability of the system being in state $\mathbf{v}$ is found by $P(\mathbf{v}) = \frac{e^{-E(\mathbf{v})}}{\sum_{\mathbf{u}} e^{-E(\mathbf{u})}}$, where $\mathbf{u}$ could be any possible state of the system. Thus, the system is most likely to be found in states that have the lowest associated energy. Restricted Boltzmann machines (RBMs) are BMs constrained to two fully-connected non-recurrent layers called the *visible layer*, where salient inputs clamp nodes to output levels of either zero or one, and the *hidden layer*, where associations between input vectors are learned. By enforcing the conditional independence of the visible and hidden layers, unbiased samples from the input can be obtained in one time-step, which enhances the learning process.

The most widely used method for training RBMs is contrastive divergence (CD), which is an approximate gradient descent procedure using Gibbs sampling [8]. CD operates in three phases: (1) *Positive Phase*: A training input vector, $\mathbf{v}$, is applied to the visible layer by clamping the nodes to either 1 or 0 levels, and the hidden layer is sampled, $\mathbf{h}$. (2) *Negative Phase*: by clamping the hidden layer to $\mathbf{h}$, the reconstructed input layer is sampled, $\mathbf{v}'$. Then, clamp the visible layer to $\mathbf{v}'$ and sample the hidden layer to obtain $\mathbf{h}'$. (3) Update the weights according to $\Delta W = \eta(\mathbf{v}\mathbf{h}^T - \mathbf{v}'\mathbf{h}'^T)$, where η is the learning rate and W is the weight matrix.

DBNs are realized when additional hidden layers are stacked on top of an RBM, and can be trained in a very similar way to RBMs. Essentially, training a DBN involves performing CD on the visible layer and the first hidden layer for as many steps as desired, then fixing those weights and moving up a hierarchy as follows. The first hidden layer is now viewed as a visible layer, while the second hidden layer acts as a hidden layer with respect to the CD procedure identified above. Next, another set of CD steps are performed, and then the process is repeated for each additional layer of the DBN.

3 SPIN-BASED BUILDING BLOCK FOR RBM

In this section, we provide a detailed description of the p-bit that provides the building block for our proposed spin-based BM architecture. Individual building blocks are interconnected by networks of memristive devices whose resistances can be programmed to provide the desired weights. For instance, in this paper, we will assume that the memristive devices are implemented using the three terminal spin-orbit torque (SOT)-driven domain wall motion (DWM) device proposed in [18].

The activation function is achieved by a spintronic building block that has been used in the design of probabilistic spin logic devices (p-bits) for a wide variety of Boolean and non-Boolean problems [3, 6, 10, 22]. The basic functionality of the p-bit shown in Fig. 1 [6] is to produce a stochastic output whose steady-state probability is modulated by an input current to generate a sigmoidal activation function. For instance, a high positive input current produces a stochastic output with a high probability of “0”, and vice versa. In

Figure 1: Structure of a p-bit.

Figure 2: Time-averaged results over 100 ns for p-bit.

the absence of any input current, the device generates either 0 or VDD outputs with roughly equal probability of 0.5, as shown in Fig. 2. This device consists of a 3-terminal, spin-Hall driven MTJ [14] that uses a circular, unstable nanomagnet ($\Delta \ll 40kT$), whereby its output is amplified by CMOS inverters as shown in Fig. 1. This MTJ with an unstable free layer can be fabricated using standard technology such that the surface anisotropy to achieve perpendicular magnetic anisotropy (PMA) that is not strong enough to overcome the demagnetizing field. Thus, the magnetization stochastically rotates in the plane, due to the presence of thermal fluctuations.

The charge current that is injected to the spin-Hall layer creates a spin-current flowing into the circular FM (in the +y direction), which does not have an axis with any preferential geometry. The spin-polarization of this spin-current is in the ($\pm z$) direction, and pins the magnetization in the (+z) or (-z) direction depending on the direction of the charge current, through the spin-torque mechanism [22]. The inherent physics of the spin-current driven low-barrier nanomagnet provides a natural sigmoidal function when a long time average of magnetization is taken. Through the tunneling magnetoresistance effect, a charge current flowing through the MTJ with a stable fixed layer detects the modulated magnetization as a voltage change. To achieve this, a small read voltage V_R is applied between the $V+$ and $V-$ terminals through a reference resistance R_0 , adjusted to the average conductance of the MTJ ($R_0^{-1} = GP + GAP/2$) where GP and GAP represented conductance in parallel (P) and anti-parallel (AP) states, respectively. This voltage becomes an input to the CMOS inverters that are biased at the middle point of their DC operating point, creating a stochastic output whose probability can be tuned by the input charge current.

Each component of the device is represented by an independent spin-circuit based on experimentally-benchmarked models that have been established in [7] and simulated as a spin-circuit in a SPICE-like platform. Here, we obtain an analytical approximation to the time-averaged behavior of the output characteristics. We

Figure 3: Proposed 32×32 hybrid CMOS/spin-based weighted array structure for RBM implementation.

Table 1: Parameters for p-bit Based Activation Function.

Parameter	Description	Value
Circular FM		
ϕ	Diameter	100nm
t	Thickness	2nm
α	Damping coefficient	0.01
MTJ		
G_0	Conductance	$150e^{-6}S$
P	Spin Polarization	0.52
Giant Spin Hall Layer(GSHE)		
λ	Spin-diffusion length	2.1nm
θ	Spin Hall Angle	0.5
Volume	$l \times w \times t$	$100nm \times 100nm \times 3.15nm$

start by relating the charge current flowing in the spin Hall layer to the spin-current absorbed by the magnet, assuming short-circuit conditions for simplicity, i.e. 100% spin absorption by the FM:

$$I_s/I_c = \beta = \frac{L}{t}(\theta)(1 - \text{sech}(\frac{t}{\lambda})) \quad (3)$$

where I_s is the spin-current, I_c is the charge current, θ is the spin-Hall angle, L , t , λ are the length, thickness and spin diffusion lengths for the spin-Hall layer. The length and width of the GSHE layer are assumed to be the same as the circular nanomagnet. With a suitable choice of the L and t , the spin-current generated can be greater in magnitude than the charge current generating “gain.” For the parameters used in this paper, which are listed in Table 1, the gain factor β is ~ 10 . Next, we approximate the behavior of the magnetization as a function of an input spin-current, polarized in the $(\pm z)$ direction. For a magnet with only a PMA in the $\pm z$ direction, a distribution function at steady state can be written analytically as below, as long as the spin-current is also fully in the $\pm z$ direction:

$$\rho(m_z) = \frac{1}{Z} \exp(\Delta m_z^2 + 2i_s m_z) \quad (4)$$

where Z is a normalization constant, m_z is the magnetization along $+z$, z is the thermal barrier of the nanomagnet, and i_s is a normalization quantity for the spin-current such that $i_s = I_s/(4q/\hbar kT)$, α being the damping coefficient of the magnet, q the electron charge and

$\hbar$ the reduced Planck constant. It is possible to use (4) to obtain an average magnetization $\langle m_z \rangle = \int_{-1}^{+1} dm_z m_z \rho(m_z) / \int_{-1}^{+1} dm_z \rho(m_z)$. Assuming $\Delta \ll kT$, $\langle m_z \rangle$ can be evaluated to give the Langevin function, $\langle m_z \rangle = L(i_s)$ where $L(x) = \frac{1}{x} - \coth \frac{1}{x}$, which is an exact description for the average magnetization in the presence of a z -directed spin-current for a low barrier PMA magnet.

In the present case, however, the nanomagnet has a circular shape with a strong in-plane anisotropy and no simple analytical formula can be derived, thus We use the Langevin function with a fitting parameter that adjusts the normalization current by a factor η , so that the modified normalization constant becomes $(4q/\hbar kT)(\eta)$. This factor increases with elevating the shape anisotropy ($H_d \sim 4\pi M_s$) and becomes exactly one when there is no shape anisotropy. Once the magnetization and charge currents are related, we can approximate the output probability of the CMOS inverters by a phenomenological equation along with fitting parameter χ as follows, $p = \frac{V_{OUT}}{V_{DD}} \approx \frac{1}{2}[1 - \tanh(\chi \langle m_z \rangle)]$, which allows us to relate the input charge current to the output probability, with physical parameters. Fig. 2 shows the comparison of the full SPICE-model with respect to aforementioned equations showing good agreement with two fitting parameters η and χ , which fit the magnetization and CMOS components, respectively.

4 PROPOSED WEIGHTED ARRAY DESIGN

Figure 3 shows the structure of the weighted array proposed herein to implement the RBM architecture including the SOT-DWM based weighted connections and biases, as well as the p-bit based activation functions. Transmission gates (TGs) are utilized in write circuits within the bit cells of the weighted connection to adjust weights by moving the DW position. As investigated in [24], TGs can provide energy-efficient and symmetric switching operation for SOT-based devices, which are desirable during the training phase. Table 2 lists the required signaling for controlling the training and read operations in the weighted array structure. Herein, a chain of inverters are considered to drive signal lines, in which each successive inverter is twice as large as the previous one.

During the read operation, write word line (WWL) is connected to ground (GND) and the source line (SL) is in high impedance

Table 2: Signaling to Control The Array Operations.

Operation	WWL	RWL	BL	SL	V+	V-
Increase Weight	VPULSE	GND	VDD	GND	Hi-Z	Hi-Z
Decrease Weight	VPULSE	GND	GND	GND	Hi-Z	Hi-Z
Read	GND	VDD	VIN	Hi-Z	VDD	VDD/2

Table 3: Relation between the input currents of activation functions and array size for $R_p = 1M\Omega$.

Features	Array Size			
	8×8	16×16	32×32	64×64
Max. Positive Current (μA)	2.71	5.14	10.79	21.46
Max. Negative Current (μA)	3.57	7.14	14.23	28.28
Max. output "0" Probability	0.77	0.88	0.95	0.97
Min. output "0" Probability	0.175	0.08	0.038	0.026

(Hi-Z) state, which disconnects the write path. The read word line (RWL) for each row is connected to VDD, which turns ON the read transistors in the weighted connection bit cell. The bit line (BL) will be connected to the input signal (VIN), which results in producing a current that affects the output probability of the p-bit device. The direction of the generated current relies on the VIN signal. In particular, since V- is supplied by a voltage source equal to VDD/2, if VIN is connected to VDD the injected current to the p-bit based activation function will have positive value, and if VIN is zero the input current will be negative. The amplitude of the generated current depends on the resistance of the weighted connection which is defined by the position of the DW in the SOT-DWM device.

During the training operation, the RWL is connected to GND, which turns OFF the read transistors and disconnects the read path. The WWL is connected to an input pulse (VPULSE) signal which activates the write path for a short period of time. The duration of the VPULSE should be designed in a manner such that it can provide the desired learning rate, η , to the training circuit. For instance, a high VPULSE duration results in a significant change in the DW position in each training iteration, which effectively reduces the number of different resistive states that can be realized by the SOT-DWM device. Resistance of the weighted connections can be adjusted by the BL and SL signals, as listed in Table 2. A higher resistance leads to a smaller current injected to the p-bit device. Therefore, the input signal connected to the weighted connection will have lower impact on the output probability of the p-bit device, which means the input signal exhibits a lower weight. The bias nodes can also be adjusted similar to the weighted connection.

5 SIMULATION RESULTS AND DISCUSSION

To analyze the RBM implementation using the proposed p-bit device and the weighted array structure, we have utilized a hierarchical simulation framework including circuit-level and application-level simulations. In circuit level simulation, the behavioral models of the p-bit and SOT-DWM devices were leveraged in SPICE circuit simulations using 20nm CMOS technology with 0.9V nominal voltage to validate the functionality of the designed weighted array circuit. In application-level simulation, the results obtained from device-level and circuit-level simulations are used to implement a DBN architecture and analyze its behavior in MATLAB.

Figure 4: Weighted array power consumption versus the resistance of the weighted connections and array size.**Table 4: Relation between the input currents of activation functions and R_p in a 32×32 array.**

Features	$R_p (M\Omega)$			
	0.25	0.5	0.75	1
Max. Positive Current (μA)	36.56	20.02	13.97	10.79
Max. Negative Current (μA)	54.95	28.12	18.9	14.23
Max. output "0" Probability	0.98	0.965	0.96	0.95
Min. output "0" Probability	0.01	0.026	0.032	0.038

5.1 Circuit-level simulation

The device-level simulations shown in Fig. 2 verified a sigmoidal relation between the input current of the p-bit based activation function and its output probability. The shape of the activation on function is one of the major factors affecting the performance of the RBM. Therefore, we have provided comprehensive analyses on the impacts of weighted connection resistance and weighted array dimensions on the input currents of the p-bit based activation functions, and the power consumption of the weighted array.

Table 3 lists the range of the activation function input currents for various weighted array dimensions, while the resistance of the SOT-DWM device in parallel state (R_p) is constant and equals $1M\Omega$. The experimental results provided in [19, 28] exhibit that an MTJ resistance in the $M\Omega$ range can be obtained by increasing the oxide thickness in an MTJ structure. The highest positive and negative currents can be achieved while the weighted connections are in parallel state, i.e. lowest resistance, and all of the input voltages (VIN) are equal to VDD and GND, respectively. The difference between the amplitude of positive and negative currents in a given array size with constant R_p is caused by the different pull-down and pull-up strengths in NMOS read transistors. The maximum and minimum output-level "0" probabilities are listed in Table 3, which can be obtained according to the measured input currents and the sigmoidal activation function shown in Fig. 2.

Moreover, Table 4 illustrates the relation between the R_p values and input currents of the activation functions, and their corresponding output probabilities, for a given 32×32 weighted array. The lower R_p resistance and higher array size provides a wider range of output probabilities which can increase the RBM performance. However, this is achieved at the cost of higher area and power consumption. The trade-offs between the array size, weighted connection resistance, and average power consumption in a single read operation is shown in Fig. 4. The lowest power consumption of

Table 5: Comparison between various RBM implementations with an emphasis on activation function structure.

Design	[12]	[15]	[2]	[21]	[4]	[9]	Proposed Herein
Weighted Connection	Embedded multipliers	Embedded multipliers	- LFSR - AND/OR gates	RRAM memristor	RRAM	PCM	SOT-DWM
Activation Function	CMOS-based LUTs	-2-kB BRAM - Picewise Linear Interpolator - Random number Generator	- LFSR - Bit-wise AND - tree adder - FSM-based tanh unit	Off-chip	- 64 × 16 LUTs - Pseudo Random Number Generator - Comparator	Off-chip	- near-zero energy barrier probabilistic spin logic device
Energy per neuron	N/A	~ 10 – 100nJ	~ 10 – 100pJ	N/A	~ 1 – 10nJ	N/A	~ 1 – 10fJ
Normalized area per neuron	N/A	~ 3000×	~ 90×	N/A	~ 1250×	N/A	~ 1×

22.6 μW is realized by an 8×8 array with $R_p = 1M\Omega$. However, this array provides the narrowest range of the output probabilities, which significantly reduces the performance of the DBN.

5.2 Application-level simulation

In the application-level simulation, we have leveraged the obtained device and circuit behavioral models to simulate a DBN architecture for digit recognition. In particular, learning rate and the shape of the sigmoid activation function is extracted by the SOT-DWM and p-bit device-level simulations, respectively, while the circuit-level simulations defines the range of the output probabilities. To evaluate the performance of the system, we have modified a MATLAB implementation of DBN by Tanaka and Okutomi [23] and used the MNIST data set [13] including 60,000 and 10,000 sample images with 28×28 pixels for training and testing operations, respectively. We have used Error rate (ERR) metric to evaluate the performance of the DBN, as expressed by $ERR = N_F/N$, where, N is the number of input data, N_F is the number of false inference [23].

The simplest model of the DBN that can be implemented for MNIST digit recognition consists 784 nodes in visible layer to handle 28×28 pixels of the input images, and 10 nodes in hidden layer representing the output classes. Fig. 5 shows the relation between the performance of various DBN topologies, and the number of input training samples ranging from 100 to 5,000, which is obtained using 1,000 test samples. The ERR and RMSE metrics can be improved by enlarging the DBN structure through increasing the number of hidden layers, as well as the number of nodes in each layer. This improvement is realized at the cost of larger area and power consumptions. Increasing the input training samples can improve the DBN performance as well, however it will quickly converge due to the limited weight values that can be provided by SOT-DWM based weighted connections. As shown in Fig. 5, some random behaviors are observed for networks with smaller sizes that are trained by lower number of training samples, which will be significantly reduced by increasing the number of training samples.

The simulation results exhibit the highest error rate of 36.8% for a 784×10 DBN that is trained by 100 training samples. Meanwhile, the lowest error rate of 3.7% was achieved using a $784 \times 800 \times 800 \times 10$ DBN trained by 5,000 input training samples. This illustrates that the recognition error rate can be decreased by increasing the number of hidden layers, and training samples, which is also realized at the cost of higher area and power overheads.

Figure 5: ERR for various DBN topologies.

5.3 Discussion

Table 5 lists previous hardware-based RBM implementations, which have aimed to overcome software limitations by utilizing FPGAs [12, 15], stochastic CMOS [2], and hybrid memristor-CMOS designs [4, 9, 21]. FPGA implementations demonstrated RBM speedups of 25-145 over software implementations [12, 15], but had significant constraints such as only realizing a single 128×128 RBM per FPGA chip, routing congestion, and clock frequencies limited to 100MHz [15]. The stochastic CMOS-based RBM implementation proposed in [2] leveraged the low-complexity of stochastic CMOS arithmetic to save area and power. However, the need for extremely long bit-stream lengths negate energy savings and lead to very long latencies. Additionally, a significant amount of Linear Feedback Shift Registers (LFSRs) were required to produce the uncorrelated input and weight bit-streams. In both the FPGA and stochastic CMOS designs, improvements were achieved by implementing parallel Boolean circuits such as multipliers and pseudo-random number generators for probabilistic behavior, which has significant area and energy overheads compared to leveraging the physical behaviors of emerging devices to perform the computation intrinsically. Bojnordi et al. [4] leveraged resistive RAM (RRAM) devices to implement efficient matrix multiplication for weighted products within Boltzmann machine applications, and demonstrated significant speedup of up to 100-fold over single-threaded cores and energy savings of over 10-fold. Similarly, Sheri et al. [21] and Eryilmaz et al. [9] utilized RRAM and PCM devices to implement matrix multiplication, while the corresponding activation function circuitry is still based on the CMOS technology, which suffers from the aforementioned area and power consumption overheads.

While most of the previous hybrid Memristor/CMOS designs focus on improving the performance of weighted connections, the work presented herein overcomes many of the preceding challenges of generating sigmoidal probabilistic activation functions by utilizing a novel p-bit device that leverages intrinsic thermal noise within low energy barrier nanomagnets to provide a natural building block for RBMs within a compact and low-energy package. As listed in Table V, the proposed design can achieve approximately three orders of magnitude improvement in term of energy consumption compared to the most energy-efficient designs, while realizing at least 90X device count reduction for considerable area savings. Note that these calculations do not take into account the weighted connections, since the main focus of this paper is on the activation function. While SOT-DWM devices are utilized herein for the weighted connections, any other memristive devices could be utilized without loss of generality.

6 CONCLUSION

Herein, we developed a hybrid CMOS/spin-based DBN implementation using p-bit based activation functions modeled to produce a probabilistic output that can be modulated by an input current. The device-level simulations exhibited a sigmoid relation between the input currents and output probability. The SPICE model of the p-bit is used to design a weighted array structure to implement RBM. The circuit simulations showed that the performance of the array can be improved by enlarging the array size, as well as reducing the resistance of the weighted connections. However, these improvements are achieved at the cost of increased area and power consumption. For instance, the lowest power dissipation among the examined designs belongs to an 8×8 array with the maximum resistance of $1M\Omega$ for weighted connections. However, this structure can only provide the output probabilities ranging from 0.175 to 0.77, which is the narrowest range among the examined designs resulting in a DBN implementation with lowest accuracy.

Next, we simulated a DBN for digit recognition application in MATLAB using the device and circuit-level behavioral models. Trade-offs include the relations between the recognition accuracy of the DBN and the number of training samples, which are comparable to conventional hardware implementations. The recognition error rate decreased substantially for the first thousand training samples, regardless of the size of the array, while benefits continue through several thousand inputs. However, at least two hidden layers are desirable to achieve suitable error rates. Finally, we have provided a comparison between previous hardware-based RBM implementations and our design with an emphasis on the probabilistic activation function within the neuron structure. The results exhibited that the p-bit based activation function can achieve roughly three orders of magnitude energy improvement, while realizing at least 90X reduction in terms of device count, compared to the previous most energy-efficient designs. The research directions herein enable several intriguing possibilities for future work, including: (1) implementing the entire network in SPICE to obtain more robust results; 2) investigating the effect of process variation and noise on the accuracy of proposed architecture; 3) studying alternative devices with lower susceptibility to thermal noise; and 4) studying the scalability challenges of DBNs using larger datasets, e.g. CIFAR.

7 ACKNOWLEDGEMENTS

This work was supported in part by the Center for Probabilistic Spin Logic for Low-Energy Boolean and Non-Boolean Computing (CAPSL), one of the Nanoelectronic Computing Research (nCORE) Centers as task 2759.006, a Semiconductor Research Corporation (SRC) program sponsored by the NSF through CCF 1739635.

REFERENCES

- [1] David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. 1985. A learning algorithm for Boltzmann machines. *Cognitive science* 9, 1 (1985), 147–169.
- [2] A. Ardakani, F. Leduc-Primeau, N. Onizawa, T. Hanyu, and W. J. Gross. 2017. VLSI Implementation of Deep Neural Network Using Integral Stochastic Computing. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25, 10 (2017).
- [3] Behtash Behin-Aein, Vinh Diep, and Supriyo Datta. 2016. A building block for hardware belief networks. *Scientific reports* 6 (2016).
- [4] M. N. Bojnordi and E. Ipek. 2016. Memristive Boltzmann machine: A hardware accelerator for combinatorial optimization and deep learning. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [5] Lars Buesing, Johannes Bill, Bernhard Nessler, and Wolfgang Maass. 2011. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS computational biology* 7, 11 (2011), e1002211.
- [6] Kerem Yunus Camsari, Rafatul Faria, Brian M. Sutton, and Supriyo Datta. 2017. Stochastic p-Bits for Invertible Logic. *Phys. Rev. X* 7 (Jul 2017), 031014. Issue 3.
- [7] Kerem Yunus Camsari, Samiran Ganguly, and Supriyo Datta. 2015. Modular approach to spintronics. *Scientific reports* 5 (2015).
- [8] Miguel A Carreira-Perpinan and Geoffrey E Hinton. 2005. On contrastive divergence learning. In *Aistats*, Vol. 10. 33–40.
- [9] S. B. Eryilmaz, E. Neftci, S. Joshi, S. Kim, M. BrightSky, H. L. Lung, C. Lam, G. Cauwenberghs, and H. S. P. Wong. 2016. Training a Probabilistic Graphical Model With Resistive Switching Electronic Synapses. *IEEE Transactions on Electron Devices* 63, 12 (Dec 2016), 5004–5011.
- [10] R. Faria, K. Y. Camsari, and S. Datta. 2017. Low-Barrier Nanomagnets as p-Bits for Spin Logic. *IEEE Magnetics Letters* 8 (2017), 1–5.
- [11] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18, 7 (2006), 1527–1554.
- [12] S. K. Kim, P. L. McMahon, and K. Olukotun. 2010. A Large-Scale Architecture for Restricted Boltzmann Machines. In *2010 18th IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*. 201–208.
- [13] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (Nov 1998), 2278–2324.
- [14] Luqiao Liu, Chi-Feng Pai, Y. Li, H. W. Tseng, D. C. Ralph, and R. A. Buhrman. 2012. Spin-Torque Switching with the Giant Spin Hall Effect of Tantalum. *Science* 336, 6081 (2012), 555–558.
- [15] D. L. Ly and P. Chow. 2010. High-Performance Reconfigurable Hardware Architecture for Restricted Boltzmann Machines. *IEEE Transactions on Neural Networks* 21, 11 (Nov 2010), 1780–1792.
- [16] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 6197 (2014), 668–673.
- [17] Ruhi Sarikaya, Geoffrey E. Hinton, and Anoop Deoras. 2014. Application of Deep Belief Networks for Natural Language Understanding. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 22, 4 (April 2014), 778–784.
- [18] Abhronil Sengupta, Aparajita Banerjee, and Kaushik Roy. 2016. Hybrid Spintronic-CMOS Spiking Neural Network with On-Chip Learning: Devices, Circuits, and Systems. *Phys. Rev. Applied* 6 (Dec 2016), 064003. Issue 6.
- [19] Abhronil Sengupta, Priyadarshini Panda, Parami Wijesinghe, Yuseung Kim, and Kaushik Roy. 2016. Magnetic tunnel junction mimics stochastic cortical spiking neurons. *Scientific reports* 6 (2016), 30039.
- [20] A. Sengupta, M. Parsa, B. Han, and K. Roy. 2016. Probabilistic Deep Spiking Neural Systems Enabled by Magnetic Tunnel Junction. *IEEE Transactions on Electron Devices* 63, 7 (July 2016), 2963–2970.
- [21] Ahmad Muqem Sheri, Aasim Rafique, Witold Pedrycz, and Moongu Jeon. 2015. Contrastive divergence for memristor-based restricted Boltzmann machine. *Engineering Applications of Artificial Intelligence* 37 (2015), 336–342.
- [22] Brian Sutton, Kerem Yunus Camsari, Behtash Behin-Aein, and Supriyo Datta. 2017. Intrinsic optimization using stochastic nanomagnets. *Scientific Reports* 7 (2017).
- [23] M. Tanaka and M. Okutomi. 2014. A Novel Inference of a Restricted Boltzmann Machine. In *2014 22nd International Conference on Pattern Recognition*. 1526–1531.
- [24] R. Zand, A. Roohi, and R. F. DeMara. 2017. Energy-Efficient and Process-Variation-Resilient Write Circuit Schemes for Spin Hall Effect MRAM Device. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25, 9 (2017), 2394–2401.