Mapping Generative Models onto Networks of Digital Spiking Neurons

Bruno U. Pedroni, Srinjoy Das, John V. Arthur, Paul A. Merolla, Bryan L. Jackson, Dharmendra S. Modha, Kenneth Kreutz-Delgado, and Gert Cauwenberghs

ABSTRACT

Stochastic neural networks such as Restricted Boltzmann Machines (RBMs) have been successfully used in applications ranging from speech recognition to image classification. Inference and learning in these algorithms use a Markov Chain Monte Carlo procedure called Gibbs sampling, where a logistic function forms the kernel of this sampler. On the other side of the spectrum, neuromorphic systems have shown great promise for low-power and parallelized cognitive computing, but lack well-suited applications and automation procedures. In this work, we propose a systematic method for bridging the RBM algorithm and digital neuromorphic systems, with a generative pattern completion task as proof of concept. For this, we first propose a method of producing the Gibbs sampler using bio-inspired digital noisy integrate-and-fire neurons. Next, we describe the process of mapping generative RBMs onto the IBM TrueNorth neurosynaptic processor – a low-power digital neuromorphic VLSI substrate. Mapping these algorithms onto neuromorphic hardware presents unique challenges in network connectivity and weight and bias quantization, which, in turn, require architectural and design strategies for the physical realization. Generative performance metrics are analyzed to validate the neuromorphic requirements and to best select the neuron parameters for the model. Lastly, we describe a design automation procedure which achieves optimal resource usage, accounting for the novel hardware adaptations. This work represents the first implementation of a generative RBM on a neuromorphic VLSI substrate.

I. INTRODUCTION

Deep Learning algorithms such as Restricted Boltzmann Machines (RBMs) and Deep Belief Networks (DBNs) have been successfully used in a wide range of cognitive computing applications such as image classification [1], speech recognition [2], and motion synthesis [3]. RBMs are generative learning algorithms and are particularly useful in extracting features from unlabeled data (i.e. unsupervised learning) [4]. Structurally, an RBM is a stochastic neural network composed of 2 layers of neuron-like units: a layer of visible units *v* which are driven by the real-world data of interest and a layer of hidden units *h* which form connections to these visible units. There are no interconnections within a layer and the weights of connections between layers are symmetric. Fig. 1a exemplifies an RBM with 4 visible and 3 hidden units. The two RBM layers form a probabilistic model of the input data based on the Boltzmann distribution as defined by the following [5]:

$$p(\mathbf{v},\mathbf{h}) = \frac{e^{-E(\mathbf{v},\mathbf{h})}}{\sum_{\mathbf{v},\mathbf{h}} e^{-E(\mathbf{v},\mathbf{h})}}, \qquad (1)$$

where

$$E(\mathbf{v},\mathbf{h}) = -\mathbf{v}^T\mathbf{W}\mathbf{h} - \mathbf{b}_v^T\mathbf{v} - \mathbf{b}_h^T\mathbf{h}.$$

Here *p* denotes the Boltzmann probability distribution and *E* is an energy function of *v* and *h*, where *v* denotes the binary state (0 or 1) of the visible units and *h* represents the binary state of the hidden units. The weight between visible and hidden units is represented by *W*, while b_v and b_h represent the biases of *v* and *h*, respectively. The denominator is the sum of all possible states of visible and hidden units, also known as the partition function.

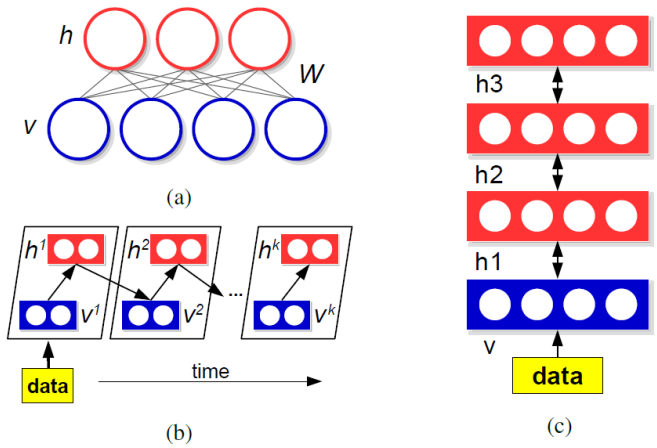

Fig. 1: RBM and DBN representations. (a) An RBM formed by 4 visible and 3 hidden units. (b) Gibbs sampling procedure in an RBM. (c) A DBN formed by stacking RBMs.

Inference in an RBM can be performed using a Markov Chain Monte Carlo (MCMC) procedure called Gibbs sampling, where each unit in any given layer is sampled based on its total input from units in the other layer. Fig. 1b illustrates *k* steps of MCMC performed in an RBM. The Gibbs sampling rule in binary RBMs is defined by the logistic function, $\sigma(x) = 1/(1 + e^{-x})$, with the probability of activation of unit *i* as defined below [5]:

$$p(x_i = 1 | x_j, j \neq i) = \sigma(\sum_j w_{ij} x_j + b_i), \qquad (2)$$

where w_{ij} is the weight from unit *j* to unit *i*, and b_i is the bias of unit *i*. DBNs are formed by stacking layers of RBMs (Fig. 1c) and it has been shown that inference in a DBN can be done in a successive layer-by-layer manner on each RBM [6]. RBMs and DBNs can be used with labeled data for classification tasks either as feature extractors to an external classifier or as completely self-contained discriminative machine learning frameworks [1]. However, most of the data in the real world is unlabeled and, in such situations, RBMs and DBNs can be used to perform generative inference tasks. Applications of inference in such unsupervised frameworks include, for example, restoration of incomplete or occluded images and prediction of motion sequences.

Currently, inference tasks using RBMs and DBNs are overwhelmingly realized in software, which are typically run on high performance CPUs (Central Processing Unit) and GPUs (Graphical Processing Unit). However, for ultra low-power, real-time realizations of these algorithms, the Neuromorphic Computing paradigm is more suitable in terms of processing and communication. Neuromorphic VLSI (Very Large Scale Integrated Circuit) devices, inspired by neural architectures and functions, have been realized with analog [7] and digital [8] circuit elements. Such systems typically compute in a massively parallel fashion and communicate asynchronously using spikes. The principal benefit of this architecture, which stands in contrast to the traditional von Neumann computing paradigm, is extremely energy efficient computation in a highly concurrent fashion. Algorithms which demand large matrix multiplications, such as RBMs and DBNs, benefit greatly in terms of computation (and, consequently, power) when implemented in spike-based systems, because multiplications by zero are avoided (i.e. absence of spike does not generate computation). Therefore, arrays of spiking neurons realized on neuromorphic VLSI are ideal for classification, generation and other inference tasks in the context of real-world high dimensional data. The goal of our work is to develop a modular architecture in a systematic fashion to form a foundation for building neural networks, such as RBMs and DBNs, on substrates of digital spiking neurons. Pattern completion performed on the TrueNorth digital neuromorphic VLSI device using the MNIST dataset – a database consisting of images of handwritten digits – serves as proof of concept of our design approach.

The remainder of this paper is divided in the following manner: Section II describes the Markov chain analysis of the digital neural sampler; Section III describes the TrueNorth system and the challenges in implementing Deep Learning algorithms, along with the necessary steps for mapping the RBM algorithm onto digital spiking neuromorphic hardware; Section IV discusses quality metrics and the impact on generative performance when using the digital neural sampler and sparse network connectivity; Section V shows the developed 3-stage RBM architecture and the generative model on TrueNorth; Section VI illustrates the spike processing flow in the TrueNorth RBM; Section VII details the design automation procedure for optimal hardware utilization; Section VIII shows how the developed architecture can be extended to other neural networks; and the last three sections discuss results, conclusions and future work, and lists the references used in this work.

II. MARKOVIAN ANALYSIS OF THE DIGITAL NEURAL LOGISTIC SAMPLER

The kernel of the MCMC procedure for inference in an RBM is the Gibbs Sampler and involves sampling from a logistic function (Eq. 2). Traditional methods for realizing a logistic sampler in hardware demand a look-up table or functional approximation [9]–[11] for the sigmoid, which is then compared to the output of a pseudo-random number generator. On the other hand, in spiking neural hardware, such as TrueNorth, the only computational primitives are neurons; therefore, it is possible to use the deterministic and stochastic neurodynamical properties for efficient realization of the logistic sampler. Below we describe the process of Gibbs sampling using digital spiking neurons in a Markov chain framework, which is useful for better understanding the sampler behavior and serves as a means for producing the generative performance metrics (detailed in Section IV).

In [12] it was shown that a digital integrate-and-fire neuron with a uniformly-sampled, integer-valued threshold combined with a Bernoulli-sampled, single-valued leak can approximate a logistic spiking probability for the corresponding RBM unit. Here we expand on this by providing a Markov chain analysis of the discrete-time neural sampler. This sampling procedure is initialized by setting the neural membrane potential (V_m) to a value equivalent to the argument of the logistic function (x in Eq. 2), and then using three neural variables (two stochastic and one deterministic) produces an approximate sigmoidal spiking probability. These variables are explained next.

1) **Stochastic leak**. The stochastic leak is an additive value to the membrane potential, and is sampled from Bernoulli trials with p=0.5. In other words, at every time step (''tick''), or the membrane potential remains the same or it is incremented by the leak value (L).

2) **Stochastic threshold**. The stochastic leak is an additive value to the base threshold value (V_{th}), and is an integer sampled from a uniform distribution. Therefore, at every tick, the membrane potential of the neuron is compared with the stochastic threshold and, in case the potential exceeds the threshold, a spike event will be generated. The stochastic component of the threshold is uniformly sampled between 0 and TR (''threshold range'').

3) **Sampling time window**. The deterministic component of the sampler is the sampling time window (T_S), which is the number of ticks during which the neuron is observed. The operation of the sampler during T_S is the following:
    a. If during T_S the neuron spikes at least once, a spike event after T_S is produced.
    b. Since the neuron may spike more than once during T_S, the sampler must still produce a single spike at the output after T_S.
    c. If no spikes occur during T_S, then no spike event should be produced at the output of the sampler after T_S.

A. ADAPTATION OF NEURAL VARIABLES INTO MARKOV CHAINS

Since we are dealing with a discrete-time digital system, the stochastic neural variables can be modeled as coupled discrete-time Markov chains (DTMC). The sampling time window determines how many steps should be taken in this chain. Each state in the chain is the instantaneous value of the membrane potential. Due to a limited number of bits for data representation, saturation levels should be taken into account. For illustrative purposes, in our examples we consider only positive leak, implying that only the positive saturation level will come into effect, as any data point beyond it will be clipped to the saturation value.

Therefore, two DTMCs with the same number of states (defined by the membrane potential range) are used to model the sampler. In terms of structure, the DTMCs will always present states representing lower-valued membrane potentials to the left of the chain, and consequently the rightmost state represents membrane potential equal to V_{sat}. Next we discuss the effect of the three neural properties on the DTMC.

1) **Stochastic leak**. Since the stochastic leak chosen for our examples causes only non-negative change in the membrane potential, the only possible transitions from a state are: (1) to itself (in the event of no leak occurrence) or (2) to the right (positive additive leak occurred). Figure 2 shows the general case of the DTMC for the stochastic leak (L). Note how no value of membrane potential can surpass V_{sat}, which makes the state representing this specific membrane potential an absorbing state [13,14]. Since it is the only absorbing state in the chain, it is called the terminating state in a terminating DTMC. Also note this state will be reached by more than one other state (not considering the self-connection) when L>1.

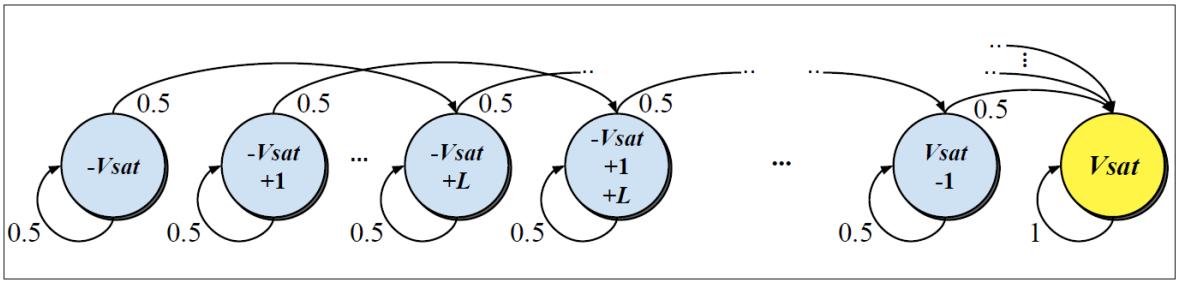

Fig. 2: DTMC for stochastic leak in neural sampler.

2) **Stochastic threshold**. The stochastic threshold is sampled from a uniform distribution, which produces a linearly increasing transition probability from states inside the range [V_{th}:V_{th}+TR] to the spiking state. Figure 3 shows the general case of the DTMC for the stochastic threshold. Note how values outside the range previously described are guaranteed not to spike ($V < V_{th}$ is realized by the self-connections) and guaranteed to spike ($V \geq (V_{th}+TR)$ is realized by the connections to V_{sat}). For simplification, in the figure the symbol $\Delta = (TR + 1)$.

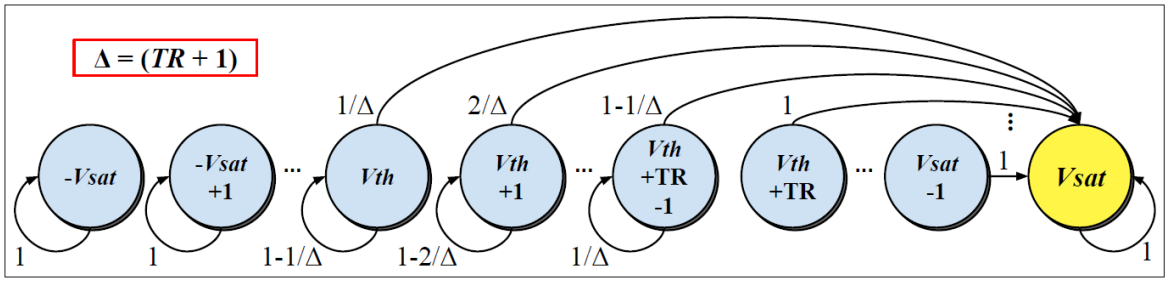

Fig. 3: DTMC for stochastic threshold in neural sampler.

To transform spikes into probabilities, we must produce a single spike event after T_S in case the neuron fired at least once during T_S. This can be obtained by using the V_{sat} state as the terminating state also for the stochastic threshold. A two-fold effect is produced by this terminating state: (1) the sigmoidal firing probability can be extracted directly from the terminating state in the stochastic threshold DTMC after T_S; and (2) the two DTMCs become coupled by using a common terminating state.

3) **Sampling time window**. The deterministic component of the sampler, the sampling time window, defines the number of steps taken in the Markov chains and is a two-phase process. The first phase occurs in the leak DTMC, where a new membrane potential value is assigned to the neuron. The second phase is the evaluation of the newly-assigned membrane potential in relation to the noisy threshold. This entire process is considered one step in the coupled DTMCs.

To transform spikes into probabilities, we must produce a single spike event after T_S in case the neuron fired at least once. This can be done by directly observing the terminating state of the combined Markov chains after T_S and, in case the neuron resides in the terminating state, V_{sat}, an ultimate single spike event will be produced; if the neuron is in any other state, no spike event will be produced.

B. MATRIX REPRESENTATION OF MARKOV CHAINS

A terminating Markov chain is a Markov chain with a single absorbing state, also known as the terminating state; all the other states are transient. The transition probability matrix -- with rows representing origin states and columns representing destination states -- of a terminating Markov chain can be defined in the following manner:

$$P = \begin{bmatrix} T & \mathbf{T}^0 \\ \mathbf{0} & 1 \end{bmatrix}$$

In matrix P, the m x m transient-states transition matrix is represented by T, the row-vector $\mathbf{0}$ represents the terminating state's non-transient transitions, and $(I_m - T)\mathbf{1} = \mathbf{T^0}$. Therefore, the entire transition matrix P can characterized by simply knowing T.

1) **Stochastic leak**. The stochastic leak is characterized by the additive leak value (L). The leak DTMC can be defined by the transition matrix P_l. The colors represent the same individual components of P_l as in matrix P.

$$P_l = \begin{array}{c} \begin{matrix} \text{-Vsat} & \text{-Vsat + 1} & \text{-Vsat + 2} & & \text{-Vsat + L} & \text{-Vsat + L + 1} & \text{-Vsat + L + 2} & & \text{Vsat - 3} & \text{Vsat - 2} & \text{Vsat - 1} & \text{Vsat} \end{matrix} \\ \begin{bmatrix} \frac{1}{2} & 0 & 0 & \cdots & \frac{1}{2} & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \cdots & 0 & \frac{1}{2} & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \cdots & 0 & 0 & \frac{1}{2} & \cdots & 0 & 0 & 0 & 0 \\ & & & & & & \vdots & & & & & \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 1 \end{bmatrix} \end{array}$$

2) **Stochastic threshold**. The stochastic threshold is characterized by the base threshold value (V_{th}) and the threshold range TR. The threshold DTMC can be defined by the transition matrix P_{th}. The colors represent the same individual components of P_{th} as in matrix P.

$$P_{th} = \begin{array}{c} \begin{matrix} \text{-Vsat} & \text{-Vsat + 1} & & \text{Vth - 1} & \text{Vth} & \text{Vth + 1} & & \text{Vth + TR - 1} & \text{Vth + TR} & & \text{Vsat - 1} & \text{Vsat} \end{matrix} \\ \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ & & & & & \vdots & & & & & & \\ 0 & 0 & \cdots & 1 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & (1-\frac{1}{\Delta}) & 0 & \cdots & 0 & 0 & \cdots & 0 & \frac{1}{\Delta} \\ 0 & 0 & \cdots & 0 & 0 & (1-\frac{2}{\Delta}) & \cdots & 0 & 0 & \cdots & 0 & \frac{2}{\Delta} \\ & & & & & \vdots & & & & & & \\ 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & \frac{1}{\Delta} & 0 & \cdots & 0 & (1-\frac{1}{\Delta}) \\ 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 1 \\ & & & & & \vdots & & & & & & \\ 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 1 \end{bmatrix} \end{array}$$

3) **Spiking probability: coupled DTMCs and sampling time window**. To obtain the sigmoidal spiking probability, the two transition matrices must first be coupled to produce $P_c = P_l P_{th}$. The spiking probability can

now be obtained by computing $P_{sample} = Pc^{T_S}$, which represents T_S steps taken in the coupled DTMC. With this, the last column (terminating state) of the final matrix will contain the spiking probability, P_{spike}, of each origin state. Therefore,

$$P_{spike}(s_i) = P_{sample}(s_i, 2V_{sat} + 1),$$

where s_i is the state corresponding to the initial membrane potential of the neuron prior to sampling.

C. EXAMPLE

The example, shown in Figure 4, illustrates the sampler obtained using the previous calculations and compared with actual stochastic neuron simulations. As can be seen, the neural sampler obtained via the coupled DTMC computation (blue line) and the stochastic simulation (10^4 samples for each initial V_m) of the neuron (blue circles) are overlapping. Besides this, the results from the DTMC computation approximate the ideal sampler scaled by a factor of 10 (red line) with considerable precision. In Section IV, a qualitative analysis of the neural versus the ideal sampler is discussed.

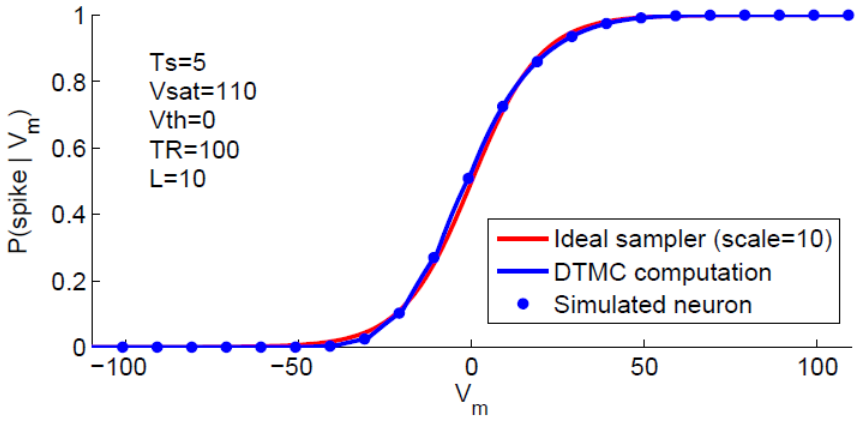

Fig. 4: Ideal sampler versus DTMC computation and neural simulation.

D. CONSIDERATIONS

The problem was analyzed for a positive stochastic leak, yet the same would be possible with a negative leak. The main detail is that the last column of P_{th}, originally considered the terminating state, would be able to be transitioned out of due to the stochastic leak in the next step of the Markov chain. Also, the terminating state in P_m would be the first column, which does not ‘’line up’’ with P_{th}.

On the other hand, if a negative leak is applied, though not sufficient for a chain in the rightmost state to reach a state below the maximum value of the threshold (i.e. below $V_{th}+TR$), then the correct spiking probability can be obtained. In this manner, even if the leak causes a transition to the left in P_m, the following iteration of P_{th} will force the system to return to the rightmost state. Interestingly, the coupled activity of the two DTMCs can preserve the original terminating state, even if it is not the terminating state in P_m.

Discrete phase-type distributions (DPTDs) are very similar in nature to the developed neural sigmoid sampler. The main difference is that DPTDs result from a system of one or more inter-related and sequentially occurring geometric distributions, while the neural sampler results from a combination of geometric (leak as Bernoulli trials) and uniform (threshold) distributions.

Lastly, the digital neural sampler is an elegant solution for sampling from a logistic function by not only using bio-inspired neural dynamics but also simultaneously realizing two operations: computing the spiking probability and sampling to obtain the new state of the unit. The DTMC computations presented will come in handy during the analysis of the generative performance of the sampler in Subsection IV-B.

III. APPROACHES FOR DEEP LEARNING ON TRUENORTH

Neuromorphic substrates present unique challenges for mapping spiking versions of machine learning algorithms on account of data precision and network connectivity constraints. In this work, the step-by-step methodology for porting RBMs and DBNs onto the IBM TrueNorth system is detailed. The MNIST dataset, consisting of 28x28 pixel grayscale images of handwritten digits 0 through 9, was chosen for both the discriminative and generative inference tasks. The following subsections present the TrueNorth system and outline the approaches and qualitative analysis of the algorithm adaptations necessary for mapping the (offline-trained) networks.

A. THE TRUENORTH DIGITAL NEUROSYNAPTIC PROCESSOR

IBM's TrueNorth is a very low-power, brain-inspired digital neurosynaptic processor [8], with 4096 cores, totaling 1 million programmable spiking neurons and 256 million configurable synapses (Fig. 5a). The core is the basic building block of the system, each composed of 256 axons (inputs) and 256 neurons (outputs) (Fig. 5b), connected via a 256 x 256 crossbar of configurable synapses (Fig. 5c). Each neuron can target its generated spikes to any axon on the chip, limited to one axon per neuron, and present over 20 individually programmable features, including threshold, leak, reset, and stochastic properties. From the user's point-of-view, neurons operate in 1 ms time steps, during which asynchronous spike event transmission and processing occurs between and inside the cores. Therefore, during each 1 ms interval, spikes are delivered to and processed in their destination cores, after which a global clock aligns the generation of the next set of spikes.

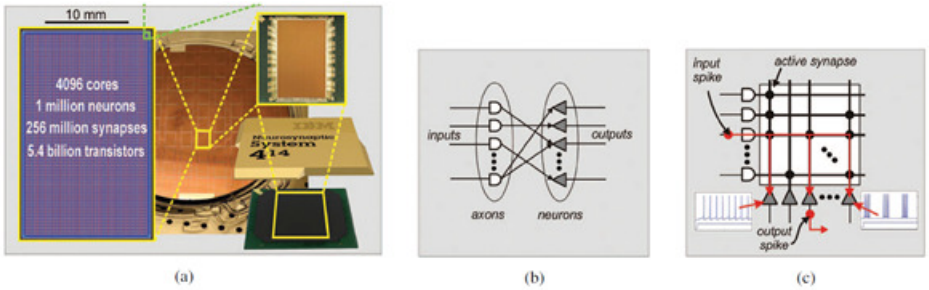

Fig. 5: TrueNorth processor and core structure [8].

The digital integrate-and-fire (I&F) TrueNorth neurons present stochastic and deterministic leak and threshold properties. A simplified representation of the dynamical behavior of the membrane potential $V_j(t)$ for neuron j at time t is defined by the following set of (sequentially processed) equations [15]:

$$\begin{cases} V_j(t) = V_j(t-1) + \Sigma_{i=0}^{255}\ A_i(t)\ w_{i,j}\ s_j^{G_i} & \text{(3a)} \\ V_j(t) = V_j(t) + (1-c_j)\lambda_j + c_j F(\lambda_j)\text{sgn}(\lambda_j) & \text{(3b)} \\ \text{if } (V_j(t) \geq \alpha_j + \eta(M_j)),\ \text{Spike and set } V_j(t) = R_j & \text{(3c)} \end{cases}$$

The first line represents the synaptic integration of all active axons at time t. The term $A_i(t)$ is the binary-valued input spike on the i^{th} axon at time t; w_{ij} is the binary-valued synaptic connection between axon i and neuron j; and s_j^{Gi} is the synaptic weight between axon i and neuron j. This last term is particularly interesting as each neuron presents four 9-bit signed integer configurable weights. Therefore, an axon can be configured to be one of four types, and this defines which of the four weights -- individually in each neuron it is connected to -- will be integrated if the axon is active.

The second line represents the leak integration, where λ_j is a 9-bit signed integer. Depending on the value of c_j, the leak can be deterministic (c=0) or stochastic (c=1). When c=0, the value of λ_i is integrated in the membrane potential. On the other hand, when c=1, the stochastic function $F(\lambda_i) = |\lambda_i| \geq \rho$ defines if a leak of sgn(λ_i) is integrated; the value of ρ is a sampled uniformly distributed 8-bit integer. The last line compares the integrated membrane potential with the threshold, which has a base value of α_j and a uniformly sampled value of $\eta(M_j)$ ranging from 0 to 2^M-1. Therefore, if $V_j(t)$ is equal to or surpasses the threshold, the neuron spikes and its membrane potential is reset to R_j.

Using the TrueNorth system as a basis for digital neural processing, the next section shows how an approximation to the Gibbs Sampler can be obtained using these neural properties.

B. GIBBS SAMPLING WITH TRUENORTH NEURONS

Neural sampling, as described in Section II, can be realized on the TrueNorth system by initializing a neuron -- which represents the RBM unit -- at a membrane potential equivalent to the argument of the logistic function. The neuron is then free to run during a sampling time window, T_S, defined in number of 1 ms time steps (''ticks''), during which a stochastic additive leak is applied and the updated membrane potential is evaluated at every tick. If the neuron's membrane potential is greater than or equal to the stochastic threshold (i.e. the neuron spikes) at least once during T_S, the binary state of the equivalent RBM unit is set to 1. Since the TrueNorth neuron may spike many times during T_S, a ''refractory'' stage was created in the RBM architecture (explained in Subsection V-D) to produce a single spike event which represents a sample from the logistic probability distribution.

The algorithm, using digital I&F neurons with stochastic leak (L) and threshold (Vth_{rand}), for realizing the sigmoidal sampling rule (equation 2) to perform MCMC sampling in RBMs is given below.

$spiked = 0$
repeat
  $V = V + \mathrm{B}(0.5)*L$
  $Vth\_rand = Vth + \mathrm{U}(0, 2^M - 1)$
  if $(V \geq Vth\_rand)$**:** $spiked = 1$
until T_S *steps*;

The term B(p) represents a Bernoulli sample (0 or 1) with probability p and U(a,b) is an integer sampled from a uniform distribution between a and b (both inclusive). In this algorithm, if the membrane potential (V) crosses the threshold (equivalent to $V \geq Vth_{rand}$), the *spiked* variable will be set to 1, after which it will remain in this state until the end of the T_S time steps. Therefore, the state of *spiked* after T_S ticks will produce an approximate sigmoidal spiking probability distribution.

For adapting the algorithm to TrueNorth, the stochastic threshold can be directly modeled by setting the appropriate values of α_r and M_r for an RBM unit represented by TrueNorth neuron r. The stochastic leak, on the other hand, cannot be directly mapped for absolute leak values greater than 1. An alternative to this is to use the additional neuron I to act as the stochastic leak for neuron r. For this, the parameters of neuron I are set to c_I=1, λ_I=+128, α_I=1, M_I=0, and R_I=0. With this, neuron n_I naturally spikes with probability $p = \lambda_I / 255 \approx 0.5$, because it leaks $\mathrm{sgn}(\lambda_I) = +1$ with this same probability and the threshold is set to α_I=1. After spiking, it is reset to V_I=0 and will present the same behavior in the next tick. If we then connect the output of neuron I to an input axon (of type i) of neuron r and set the memory position s_r^{Gi} equal to the exact leak value, we will obtain the desired spiking behavior. Lastly, multiple spikes may occur during T_S, and this is taken care of by the refractory stage and is covered in detail in Section V.

As was shown, the argument of the logistic function (x in Eq. 2) is modeled as the membrane potential of the neuron. Since TrueNorth neural membrane potential takes on only signed integer values, and the logistic function has a dynamic range between approximately -6 and +6, it is necessary to apply a multiplicative scaling factor, s, to the RBM weights and biases to increase the dynamic range of the neural logistic sampler realization. As a result of this scaling, the neural sampler must be realized with appropriate values of T_S, V_{th}, M, and L to enable the RBM to sample correctly from the logistic probability distribution. The ideal sampler (with s=50) is compared with the TrueNorth realization (T_S=8, stochastic threshold ranging from 79 to 590, and stochastic leak of 49) in Fig. 6.

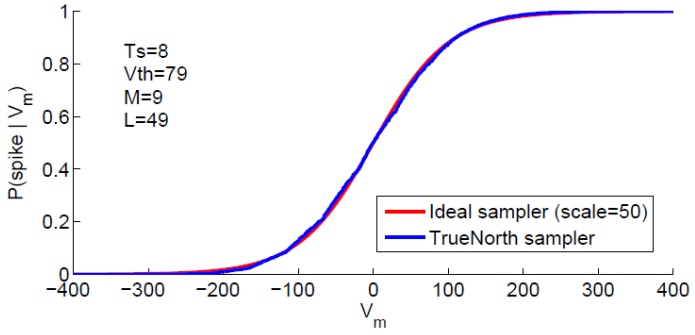

Fig. 6: Logistic sampler on TrueNorth.

C. SPARSE CONNECTIVITY

The all-to-all connectivity between layers in the RBM algorithm implementation has to be adapted to the available connectivity in hardware. The 256-input cores in TrueNorth present a constraint for the 784 pixel images used in the application, where each hidden unit ''observes'' all 784 visible units. A viable solution is to use a patching scheme over the original image, thus reducing the area of the image observed by each hidden unit. Reciprocally, the quantity of hidden units should also be selected in a way as to reduce the number of units observed by the visible units (generative task). Fig. 7a shows how a patch (yellow) is formed by an 8x8 pixel window over an MNIST image.

In [16], square patches of size p x p were randomly placed over the input image, with all the visible units belonging to a patch connected to a single hidden unit. Though this resulted in reduced network connectivity, for a physical implementation with fan-in constraints a systematic patching scheme is necessary to produce a well established maximum number of units observed in each layer. The systematic patching is particularly important for the generative model due to the feedback from hidden to visible units during inference (details in Section V), which, if patching were performed randomly, could cause a visible unit to be captured by more than 256 patches, making this fan-in unfeasible on TrueNorth. Therefore, we applied an overlapping, yet deterministic, patching scheme developed for the generative RBM realization. The method uses patches with p^2 pixels which are formed by ''sliding'' a square window over the N^2-pixel image and forming a new patch at every new position. The total number of overlapping patches produced using this method is defined by:

$$\text{patches} = \text{hidden units} = (N - p + 1)^2. \qquad (4)$$

A patching scheme implemented in this manner can be interpreted as applying a mask, W_{mask}, over the RBM's weight matrix, where 0's and 1's in the mask represent, respectively, no connection and presence of connection between visible and hidden units. The mask is applied during the offline RBM training and the resulting sparse weight matrix is then used for mapping the RBM onto TrueNorth. In Fig. 7b, the sum of column values in each row of W_{mask} represents the number of hidden units observed by each visible unit, and the sum of row values in each column represents the number of visible units observed by each hidden unit. With this systematic patching, notice that the bounds of the sums (both in rows and columns) are well-defined.

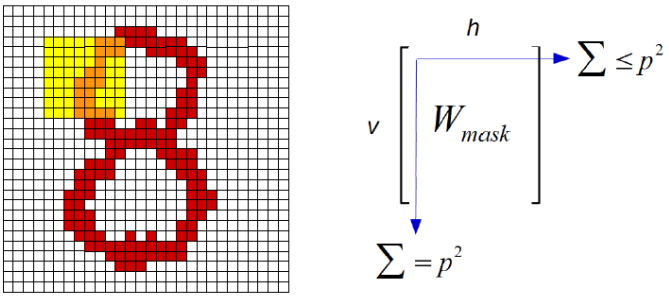

Fig. 7: Sparsity structure in RBM. (a) Illustration of an 8x8 pixel patch (in yellow). (b) Sparsity can be seen as applying a mask over the network's weight matrix.

IV. QUALITY METRICS OF DIGITAL NEURAL SAMPLER AND SPARSE NETWORK

The following subsections present qualitative analyses of the impact of the adaptations demanded during the mapping of the original RBM algorithm onto TrueNorth. First, the impact of quantization due to digital hardware data representation is verified. Second, the effect of approximate logistic sampling using the digital neural sampler is analyzed. Lastly, due to non-viability of all-to-all connections between neurons in TrueNorth, we analyze the impact of sparsity in network connectivity.

For the neuromorphic adaptations, the generative performances are verified using the Kullback-Leibler (KL) divergence and/or Annealed Importance Sampling (AIS), which are briefly explained next.

Kullback-Leibler divergence. KL divergence is a measure of the difference between probability distributions. Since we are comparing the performance of the samplers versus the exact probability distribution, only small networks with tractable partition functions (i.e. with a countable normalizing sum of all state probabilities) can be analyzed. This

measure is particularly important for our analysis of the digital neural sampler and, though we cannot directly extrapolate values of KL divergence to larger networks, the results aid in identifying expected performance for each sampler. KL divergence is defined by the following equation [17]:

$$D_{KL}(P||Q) = \sum_i P(i) log \frac{P(i)}{Q(i)}, \qquad (5)$$

where P and Q are two probability distributions, and D_{KL} is always non-negative. For our experiments, Q_i was defined to be the exact probability of the network being in state i.

Annealed Importance Sampling. AIS is a metric used to estimate the log-probability of a generative model [18,19]. For high-dimensional models, such as RBMs, where calculation of the partition function is intractable, the AIS algorithm is very useful as it performs a stochastic estimation of the partition function to compute the log probability of the model with respect to the data. Therefore, the AIS algorithm will be used for validating the generative performance of the sparsely connected network.

A. QUALITY OF DATA QUANTIZATION

The effect of data quantization can be verified using the ideal sampler. The weights and biases can be quantized by realizing the following: multiply the values by a scaling factor (s), then round the result to the nearest integer, and finally divide the second result by s. The KL divergence was computed over 1000 experiments, each consisting of randomly generated weights and biases for a network with 5 visible and 5 hidden units. For these, based on experimental results of weights and biases from previously trained RBMs, the values were sampled from the following normal distributions: weights ~ N(-0.05, 1.6e-3), visible biases ~ N(-0.3, 1), and hidden biases ~ N(0.5, 2.25). The KL divergence results are shown in Fig. 8, including a box plot of KL divergence for s = 20-100. A saturation point can be seen around s = 50.

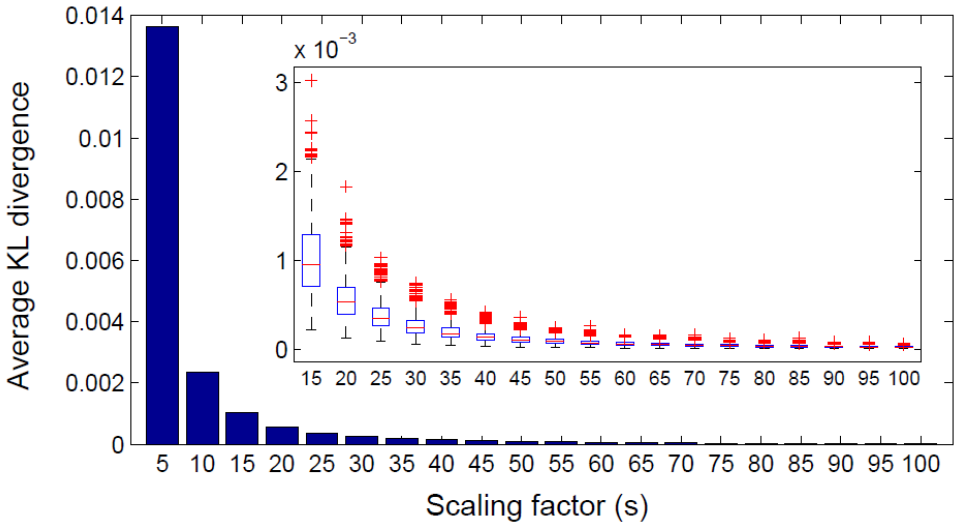

Fig. 8: Generative performance versus quantization.

The AIS results of data quantization were also verified (not shown), however they showed more variance, though an apparent saturation point was found around s = 15. On the other hand, by (subjective) visual analysis, the pattern completion task showed acceptable results only for $s \geq 30$, and therefore we based our scaling factor of s=50 for the physical implementation based on the KL divergence results. It is important to note that very large values of s are beneficial for the algorithm, however they can be costly in terms of hardware resources (cores, in the case of TrueNorth), since more neurons and longer accumulation times will be required for mapping larger values of weights and biases (explained in Section V).

B. QUALITY OF THE DIGITAL NEURAL SAMPLER

The scaling factor impacts the resource usage of the system and, as will be shown in Section VI, it also impacts the latency. The other parameter which also affects latency is the sampling time window (T_S). Thus, when neural samplers present the same generative performance, the selected configuration will naturally be the one presenting the lowest T_S. Additionally, when selecting the neuron parameters (''configuration''), one important aspect of TrueNorth neurons that should be taken into account is the membrane potential range. From Eq. 3c, we can observe that the upper bound (i.e. positive saturation) value of membrane potential is defined by the sum $\alpha_j + \eta(M_j)$. Therefore, for our analysis, we have chosen configurations with this sum close to or surpassing the upper bound of the dynamic range of the sigmoid function (≈ 6x scaling factor) while still presenting adequate sigmoid fitting.

To begin the neuron parameter selection, we first fix the scaling factor, then the quality of the digital neural sampler can be verified by varying the sampling time window (T_S) and adjusting the neuron parameters (V_{th}, M, and L) to fit the sigmoid. The best fit was found for each T_S value by performing a parameter search to reduce the mean squared error (MSE) between the ideal (scaled) logistic function and the curve produced by the neural sampler. For our experiments, five configurations were chosen, with the TrueNorth neuron parameters of the configurations (G1-G5) shown in Table I. The MSE of each neural sampler versus the ideal logistic function is presented in the rightmost column.

Config.	Scaling factor	T_S	Vth	M	L	MSE
G1	50	1	0	7	125	0.4878
G2	50	2	0	8	100	0.1311
G3	50	4	66	8	77	0.0741
G4	50	8	79	9	49	0.0412
G5	50	16	186	9	36	0.0415

Table I: Neuron configurations for neural sampler analysis.

The generative model performance for these configurations was verified by means of average KL divergence of the model versus the ideal sigmoid sampler over 10 randomly sampled networks (5 visible and 5 hidden units), with 15 experiments run for each network, and each experiment consisting of 10^5 samples. Fig. 9 shows the average KL divergence results of the different parameter configurations. The smaller plot in this figure is a boxplot of the 150 (10 networks x 15 experiments per network) KL divergence values at sample 10^5. Naturally, the configurations with lower MSE also presented lower KL divergence, with G3 and G4 practically overlapping.

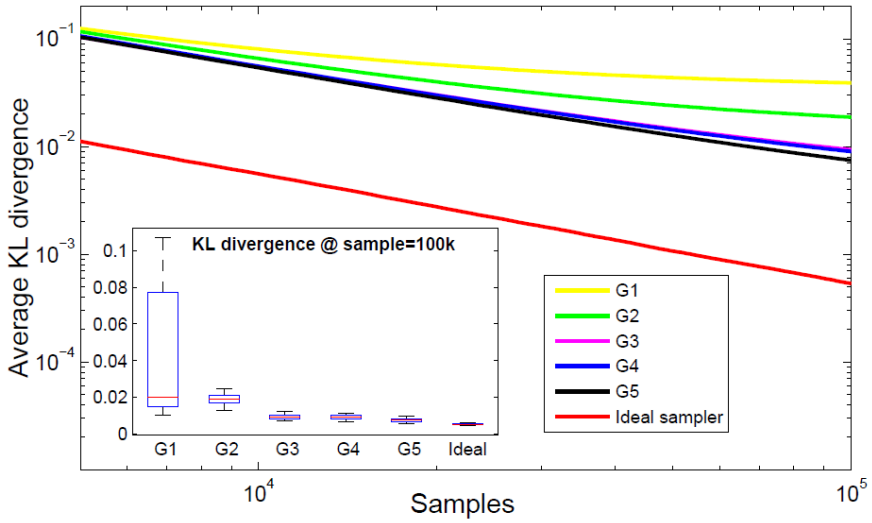

Fig. 9: Generative performance of neural Gibbs samplers.

As was mentioned in the end of Section II, the DTMC computations of the neural sampler can be very useful when simulating the network dynamics. Instead of having to simulate every step of the neuron during the sampling time window (T_S), we can simply use the spiking probability curve obtained from the DTMC as the neuron's transition operator. In other words, the probability of spiking after T_S can be extracted from the curve and this value is then compared to a uniformly-sampled number between 0 and 1.

A comparison of the normalized MSE and KL divergence (at sample 10^5) is shown in Fig. 10. Though the results for both measures were not identical -- for example, the MSE for G4 and G5 were basically identical, yet the KL divergence for G5 showed a slight improvement --, the figure clearly shows similar trends for both measures. Thus, these results indicate that using the DTMC analysis of the sampler combined with the MSE measure can be a powerful tool for quick access to estimating the generative performance of a sampler. For the generative RBM on TrueNorth, configuration G5 was chosen due to slightly better KL divergence results.

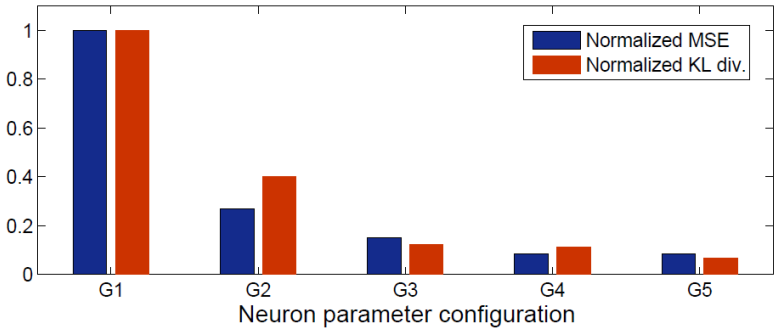

Fig. 10: MSE and mean KL divergence of neural Gibbs samplers.

C. QUALITY OF THE SPARSE NETWORK

Since sparsity is difficult to be evaluated in small networks, the generative qualities of the sparse RBM were verified by means of the AIS measure of a network with 784 visible and 500 hidden units. In [16], the authors show how a sparsely connected RBM can produce a more noise-tolerant model for classification. For our application, sparsity is actually necessary for reducing the fan-in of each neuron, which, in TrueNorth, is limited by the 256-input cores. The patching scheme proposed for the generative RBM, described in Subsection III-C, takes into account the feedback from hidden to visible units. With this method, as illustrated in Fig. 7b, the patch dimension (p) defines the maximum number of observed units in both directions (i.e., visible→hidden and hidden→visible).

AIS measure versus patch dimension results are shown in Fig. 11. For low p values, lower log probabilities were produced on account of less information captured by each patch (p^2<<). For large p values, the log probability is also lower on account of less number of hidden units (refer Eq. 4). Given the performance results of the model, for the generative RBM implementation an optimal patch size of 8 x 8 was chosen, resulting in $(N\text{-}p\text{+}1)^2 = (28\text{-}8\text{+}1)^2 = 441$ hidden units.

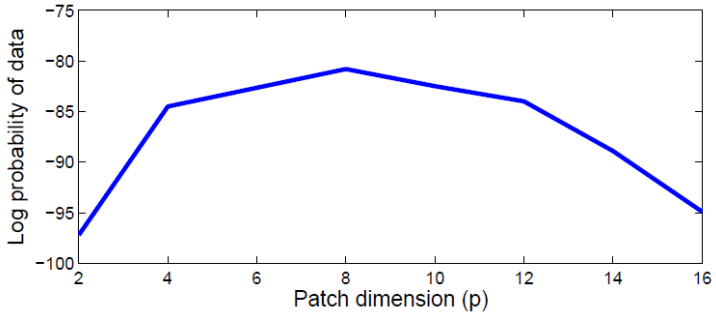

Fig. 11: Generative performance versus sparsity.

V. GENERATIVE RBM ARCHITECTURE ON TRUENORTH

The generative RBM was mapped on TrueNorth by developing a modular 3-stage architecture, where each combination of these three stages represents the transition between RBM layers. The diversity of configurable parameters present in TrueNorth is critical to the realization, with particular neuron types, connectivity strategies, and reset modes in each stage. The physical constraints of TrueNorth -- particularly 256 axons and neurons per core, only 1 destination axon per neuron, and 4 distinct weights per neuron -- defined the design flow of the RBM. The architecture, composed of stages (1) refractory-and-splitter, (2) quantization and (3) accumulate-and-sample, is illustrated in Fig. 12a.

The generative application implemented was a pattern completion task of a corrupted MNIST image. The signal flow in the TrueNorth RBM is illustrated in Fig. 12b, where each row is a 3-stage module and the blue and red blocks represent information related to visible and hidden units, respectively. Note the second stage in each module contains both colors, since this is the transition between visible and hidden layers, i.e. where the arguments of $p(v|\mathbf{h})$ and $p(h|\mathbf{v})$ are computed. Finally, the data flow for the application is represented in Fig. 12c. For this task, part of an image of the digit ''6'' (not used during training) was removed, and the figure shows the first 3 reconstructions based on the partial data.

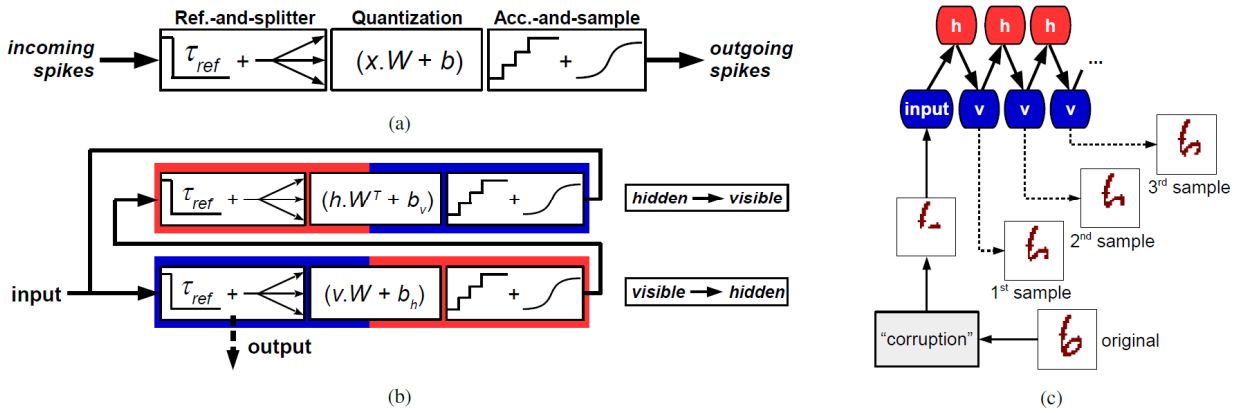

Fig 12.: TrueNorth RBM. (a) The 3-stage architecture used to distribute spike events (splitter), produce the desired membrane potential, and realize the sigmoid sampling. (b) The generative model structure, formed by combining two 3-stage blocks and including the feedback between layers. (c) Example of a pattern completion task where the digit ''6'' is incrementally reconstructed.

A. STAGE 1A: SPLITTER

Stage 1 serves a dual role in the system: (1) a splitter for input signals in the respective RBM layer and (2) a refractory effect of the neuron. The splitter is necessary because of the one-to-many connections from one layer of the RBM to the other, and therefore stage 1 generates the required number of replicas of an RBM unit to be used in the quantization stage. Fig. 13a illustrates a splitter core used for generating the necessary number of replicas of each of the visible units. The neurons are set to unit thresholds and all synaptic connections are of weight equal to +1, which will cause the neurons to spike whenever an axon event arrives. The refractory function of this stage and the two control signals (C^+ and C^-) are discussed later in Subsection V-D.

B. STAGE 2: QUANTIZATION

In TrueNorth, the weight of connections between axons and neurons can be configured with two constraints: the weights between axons connected to a given neuron are allowed to have only 4 different values; and each axon is configured as one of 4 types, reflecting on which of the 4 weights will be used for the connection between the axon and the respective neuron [15]. The first constraint limits the number of different possible weights, while the second limits the ''reutilization'' of axons between neurons. This is because an axon can be used amongst two neurons only if the weight stored in each neuron's memory position -- defined by the axon type -- is the desired synaptic weight for each of these connections. Since the diversity of weights and biases demanded by the generative RBM cannot be directly represented by this memory structure, a quantization stage is necessary to realize the connectivity between RBM units.

The representation of individual RBM weights and biases was achieved by using a collection of neurons in stage 2, each with its own weight, which together can produce the desired membrane potential (i.e. the equivalent argument of $\sigma(x)$). For this, linear-reset, unit-threshold neurons are used [15], and they operate by decrementing their membrane potential by 1 every time they spike, continuing to do so while the value is above zero. In this manner, the collective activity of many stage 2 neurons encodes the RBM weight/bias, while stage 3 will be used to accumulate the spikes from these many neurons into a single neuron.

The quantization of weights and biases is done by selecting a maximum accumulation time (T_A), which will be the largest value of membrane potential a stage 2 neuron can reach. In other words, every input spike into stage 2 axons will charge the membrane potential of each quantization neuron up to at most T_A, after which they will freely operate,

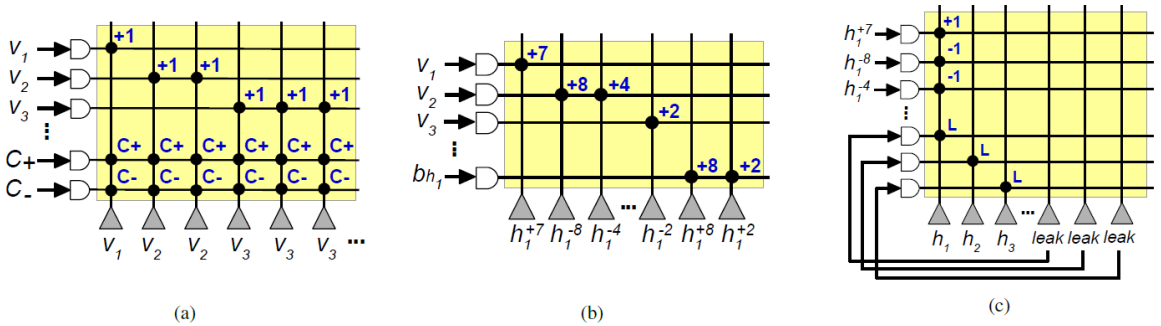

Fig. 13: Example of TrueNorth RBM stages: (a) Refractory-and-Splitter, (b) Quantization, and (c) Accumulate-and-Sample.

with spiking activity guaranteed to cease in a maximum of T_A ticks. Fig. 13b exemplifies a stage 2 core with T_A=8 and visible units v_1, v_2, and v_3 connected to hidden unit h_1 with weights +7, -12, and -2, respectively. Since T_A is a user-defined value, intuitively we would select the lowest value possible as to reduce the overall latency of the system. However, depending on the number of weights to be mapped and their specific values, attempting to use smaller values of T_A will exceed the number of available neurons in a core. Note that the sign for the negative weights is actually positive, for stage 2 only takes into account the intensity (absolute value) of the connection between units, independent of being excitatory or inhibitory. The actual sign of the connection is taken care of in stage 3. Lastly, since bias values are independent of neuronal activity, these are realized by sending an external spike event to the bias axon (b_{h1}=10 in Fig. 13b) each time the sum of inputs to a given RBM unit neuron is to be computed.

C. STAGE 3: ACCUMULATE-AND-SAMPLE

Stage 3 is used to accumulate the activity of the quantization neurons into a single neuron, which will then be sampled (as described in Subsection III-B). Prior to accumulation, the membrane potentials of the stage 3 neurons are initialized to zero. Then, during the first time window (T_A), stage 3 neurons accumulate spikes from stage 2 neurons to form a membrane potential equivalent to the argument of the logistic function. The neurons used in stage 3 have a non-resetting property to prevent clearing the membrane potential during the accumulation phase. This is followed by the time window T_S, during which the stochastic threshold and leak properties of the neuron are used for sampling from the logistic probability distribution. During this second time window, if the neuron's membrane potential surpasses the threshold, the neuron may spike multiple times since it is configured as non-resetting. For the spikes to correctly represent a sample from the logistic function, the refractory stage is necessary to register a maximum of 1 spike event per sampling window, and is described in the next subsection.

An example of a stage 3 core crossbar configuration is shown in Fig. 13c, with the sign of the RBM weight/bias now included in the synaptic weights. Note the stochastic leak is realized using a recurrent connection from a dedicated neuron, which samples from a Bernoulli distribution with p=0.5. This is necessary because the absolute value of the stochastic leak in TrueNorth is limited to 1, so an additional synaptic weight was used to produce larger values of L for sampling the neurons as required by the digital neural sampler.

D. STAGE 1B: REFRACTORY

Spike events generated by stage 3 neurons of an RBM layer are sent to the stage 1 axons of the other layer. To convert the multiple spikes received by stage 1 axons into a single spike, and align all the events in this stage at the same 1 ms time step, the splitter neurons were configured with a negative saturating membrane potential. The membrane potential of stage 1 neurons are initialized to the negative saturating value C^-, with $|C^-| > T_S$, at the start of the sampling phase of stage 3 in the other RBM layer. Every incoming spike in stage 1 will cause the membrane potential of its associated neuron to increase by 1. After T_S, the membrane potential of the stage 1 neurons are incremented by C^+ (= $|C^-|$), causing the neurons which received at least one spike to cross the threshold and simultaneously generate a spike to the subsequent stage 2.

VI. SPIKE PROCESSING FLOW IN TRUENORTH RBM

An example of the spiking activity flow between RBM layers is shown in Fig. 14, where the following parameters were used: T_S=10, T_A=8, C^-=-30, stage 3 stochastic threshold ranging between 10 and 17 and stochastic leak of +3. In the example, the *x*-axis denotes time (in 1 ms ticks) and the *y*-axis denotes the value of the membrane potential (V_{mem}). The blue line is the neuron's membrane potential, the solid red line is the saturation level, the dashed red line is the threshold, and the red circles represent spike events.

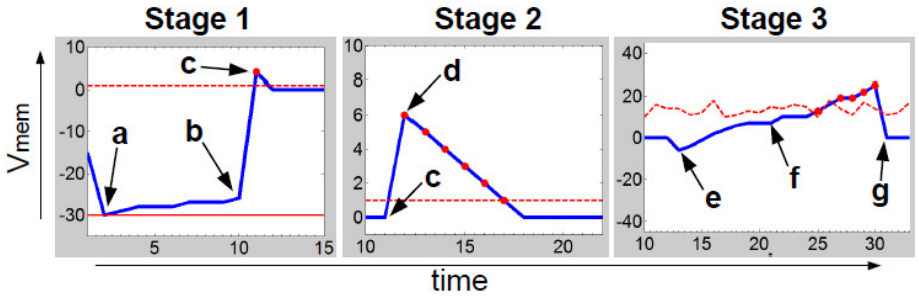

Fig. 14: Example of spike processing flow. Stage 1 realizes two functions: refractory effect in (a) and (b), and splitter (spike distribution) in (c). Stage 2 quantizes the weights between RBM layers in (c) and (d), producing the desired membrane potential for stage 3 to sample from. Stage 3 accumulates the spikes from stage 2 linear-reset neurons between (e) and (f), and the sampling procedure is performed between (f) and (g).

The sequence of events (the letters) presented in Fig. 14 are detailed below:

a. **Time=2**. Stage 1 neurons are initialized to C^-, after which they begin accumulating spikes from stage 3 neurons of the other RBM layer.
b. **Time=10**. After T_S, the C^+ signal is applied, and every neuron which captured at least one spike from stage 3 neurons crosses the threshold=1.
c. **Time=11**. As C^+ is applied, spike events from stage 1 neurons are transmitted to stage 2 axons. In this example, the stage 2 neuron is charged to a membrane potential of 6.
d. **Time=12-17**. The linear-reset stage 2 neurons continuously produce spike events to stage 3 axons until their membrane potentials return to zero.
e. **Time=13**. At this moment, the stage 3 neurons begin accumulation for T_A ticks.
f. **Time=21**. After the stage 3 neurons have accumulated their membrane potentials to the desired values, the sampling phase begins. The stage 1 neurons of the other RBM layer are initialized to C^-; the stochastic threshold and leak come into effect at stage 3.
g. **Time=31**. After T_S ticks, the stage 3 neurons are reinitialized.

VII. DESIGN AUTOMATION

TrueNorth system configuration can be realized using the object-oriented *Corelet* Language, which is an abstraction for representing the network of neurosynaptic cores [20]. The developed design automation procedure consists of creating systematic data structures, originating from the RBM weight and mask matrices, RBM biases, and user-defined parameters, which include: accumulation time T_A; sampling time (T_S); data scaling factor (s); and sampler stochastic threshold and leak. Once these have been defined, the automation procedure produces an optimal configuration of cores which minimizes the number of axons and neurons used for the RBM realization. This procedure is described below. Note that all considerations for hidden units are also valid for visible units.

Strategy 1.1: The first strategy involves establishing the number of neurons required for mapping each RBM weight and bias. Without optimization in stage 2, the number of neurons n_j used when quantizing the weight between the visible units observed by hidden unit h_j can be computed by $n_j = \sum_i \text{ceil}(w_{ij}/T_A)$. This direct method of mapping weights and biases does not take into account the fact that possibly many stage 2 neurons present low weights, which will cause them to complete spiking (during the accumulation phase) before neurons which represent higher values, such as weight T_A. Since the network must always go through T_A ticks during the accumulation phase, it would be more efficient to try to connect a given neuron to as many possible axons, provided the total synaptic weight is guaranteed not to exceed T_A. In the limiting case, neurons which map weights -1 and +1 can have up to T_A axons connected to them.

Though this first strategy benefits the core utilization considerably, better optimizations are possible. This is because the order in which the RBM weights are chosen to be mapped in stage 2 is defined by the user, yet different mapping sequences may utilize less cores. For example, suppose T_A=4 and the weights to be mapped are 1 through 6 for visible units v_1 through v_6, respectively. If we were to map them in this order, a total of 6 neurons would be used (Fig. 15a). On the other hand, if we were to map in the reverse order (6 through 1), a total of 7 neurons would be necessary (Fig. 15b). Therefore, the order of weight mapping affects the core utilization. Since the possible number of weight orderings to be analyzed is intractable, better results can be obtained by using strategy 1.2.

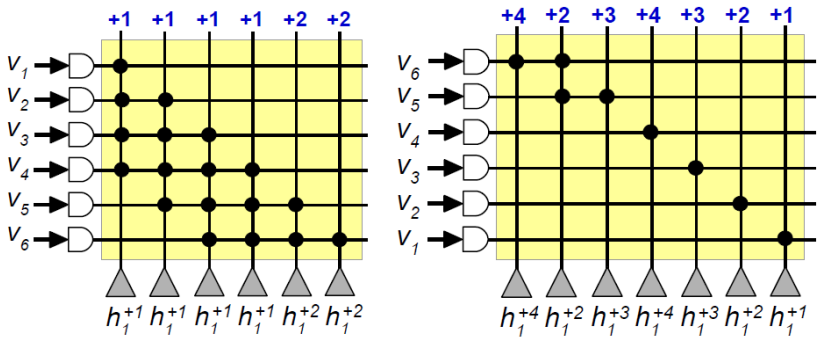

Fig. 15: Strategy 1.1 examples of stage 2 quantization.

Strategy 1.2: In this strategy, the weights closest to a user-defined central weight value are mapped first. By sweeping through all possible central weights, an optimal value can be empirically obtained. Fig. 16 shows the number of neurons used when mapping the weights -20 through 20 with T_A=4. The red line is the number of neurons (120) obtained with no optimization, while the black line is the number (110) when using sequential mapping with weight neuron ''reutilization'' (i.e. strategy 1.1). The blue line shows the results for the central weight method (strategy 1.2). The reduction from 110 to 107 neurons when using a central weight of 5, for example, is small (approximately 3%), though more significant reductions are possible when this strategy is combined with strategies 2 and 3.

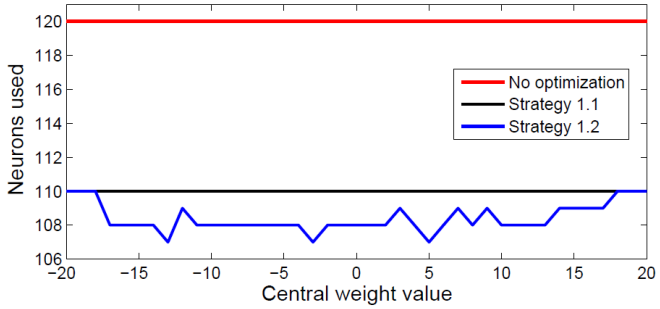

Fig. 16: Optimization strategy comparison in terms of number of neurons used to map the desired weights.

Strategy 2: One final optimization can be performed in stage 2. When a hidden unit is mapped, the number of remaining neurons in the core may be enough to map additional hidden units. If this is the case, among the hidden units to be mapped, we select the one which has the most number of visible units in common with the hidden units previously mapped in the given core. This is because the patching scheme produces hidden units which may have some visible units in common and, thus, specific units are capable of sharing axons. This strategy is naturally also valid when mapping visible units.

Strategy 3: Stages 1 and 3 can also be optimized via a greedy optimization method. For stage 1, the algorithm selects the unit which uses the most number of neurons (replicas for axons in stage 2), yet does not exceed the neuron limit in the core. If no unit can be mapped in the core, the algorithm creates a new one, until all units have been mapped. For stage 3, the algorithm does the same as for stage 1, now selecting the unit which uses the most number of axons (quantization neurons from stage 2) without exceeding the axon limit in the core.

VIII. EXTENSION TO OTHER MODELS AND NETWORKS

Using the 3-stage architecture developed for the generative RBM, discriminative models on TrueNorth can also be implemented. Previously, two other TrueNorth RBM implementations [21,22] were also developed for MNIST image classification. In [21] an RBM was implemented simply as a feature extractor, where classification was performed by sending hidden unit spikes off-chip into a linear classifier. In [22], the RBM was also used as a feature extractor, except classification was now realized using an on-chip spiking linear classifier. On the other hand, using our proposed architecture, a completely self-contained discriminative RBM could be implemented. Figure 17a shows this proposed model, with the class units now included in the RBM (in yellow).

The main difference in this architecture would be the feed-forward aspect of spike flow, as opposed to the feedback from hidden to visible demanded by the generative model. Another important aspect of the discriminative RBM is that the ''refractory'' component of the bottom module in Fig. 17b would not be required, as this component would actually be demanded at the output of the top module. However, the refractoriness of the spikes generated by the class units can be realized off-chip since basically all the computations and processing have already been realized by the TrueNorth system. As an illustrative example, for MNIST digital classification the output of the system would be the spiking activity of the class units, where Fig. 17a shows a sample raster plot of the class units for an input image of a ''5'' being correctly predicted (i.e. class unit 5 spiked the most).

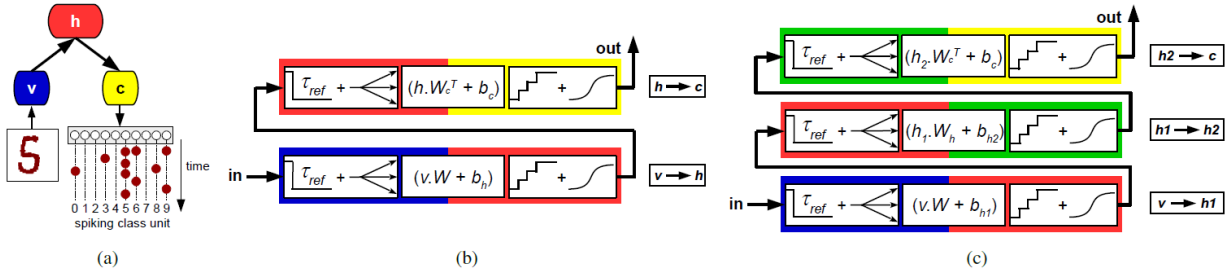

Fig. 17: Extension of 3-stage architecture to discriminative RBM (a) during inference and (b) TrueNorth architecture. (c) Discriminative Deep Belief Network (DBN) architecture with additional (green) hidden layer.

Naturally, by stacking RBMs one could extend the 3-stage architecture to deeper network structures. In Fig. 17c, a discriminative DBN with two hidden layers (in red and green) is shown. As with the previous network, the stochastic spiking activity of the class units (*out*) can be used as prediction of the currently observed data (*in*). Generative DBNs could also be implemented, where the feedback would again be required.

Lastly, due to the distinct functions of each stage in the 3-stage architecture, it is possible to adapt the stages for other types of networks. Stage 1 will be usually necessary since many applications deal with densely-connected networks. Stage 2 will be used when high-fidelity data representation is necessary -- as was the case of generative RBMs. The most differentiated stage, perhaps, is stage 3, which, for the RBM, is where the stochastic properties of the network are active and sampling occurs. Therefore, stage 3 can be adapted for deterministic applications, as well as for other sampling kernels and techniques.

IX. RESULTS

For realizing the generative model -- the MNIST pattern completion task detailed in Section V -- on TrueNorth, an RBM with 784 visible units and 441 hidden units (generated by using 8 x 8 patches) was trained using the persistent Contrastive Divergence algorithm [23]. The generative application demands a sampler with high fidelity with respect to the ideal sampler. To achieve this, the parameters were selected according to the criteria outlined in Section IV: scaling factor = 50, T_S=16, stochastic leak = 36, and stochastic threshold ranging from 186 to 697. The choice of the scaling factor directly impacts the RBM weight and bias magnitudes. To map these weights in stage 2, a trade-off is necessary between the accumulation time and the quantity of neurons and cores demanded by the application. Therefore, a compromise value of T_A=32 was selected for the mapping.

Using the automation strategies outlined in Section VII, the generative RBM was realized with 865 cores, representing 21% of the total number of cores on TrueNorth. Table II shows how applying the optimization strategies 1.2, 2 and 3 drastically reduced the core utilization.

Case	Strategies	# of cores	Chip utilization (%)
1	none	2956	72.2
2	1.1, 2, 3	906	22.1
3	1.2, 2, 3	865	21.1

Table II: Core utilization results based on optimizations.

Example outputs of the pattern completion task are shown next. In Fig. 18a, one example output for each of the ten digits is presented: the first column is the original data (''O''), the middle column is the corrupted (''C'') image sent into the TrueNorth RBM, and the third column is the reconstructed (''R'') output after 50 samples. These images were chosen to represent positive results, while Fig. 18b shows images whose reconstruction was not ideal. Lastly, Fig. 18c illustrates a sequence of reconstructions for a corrupted image of the digit ''6''; the sample number is indicated above each image.

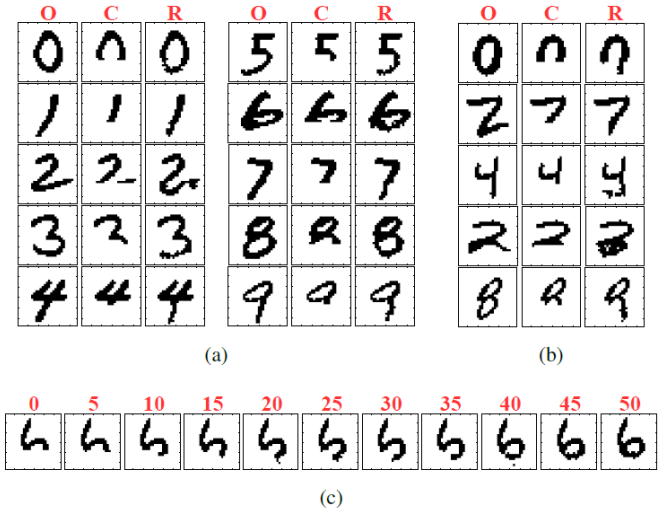

Fig. 18: TrueNorth pattern completion task outputs. Positive (a) and negative (b) reconstruction results. Reconstruction of the digit ''6'' in (c), with the sample number indicated above each image.

Depending on the percentage of image occlusion (''corruption''), the RBM may or may not be able to reconstruct a satisfactory representation of the original image. Therefore, we performed experiments with different image occlusion percentages and the Hamming distance (HD) -- which, in this case, is the same as the number of incorrectly reconstructed pixels -- was measured over 1,000 test images at the 50th reconstruction sample. The mean value of the HDs was normalized according to the number of non-occluded pixels in the image. The results, illustrated in Fig. 19, show that the reconstructive performance of the neural sampler nearly matches that of the ideal sigmoid sampler (in terms of normalized Hamming distance).

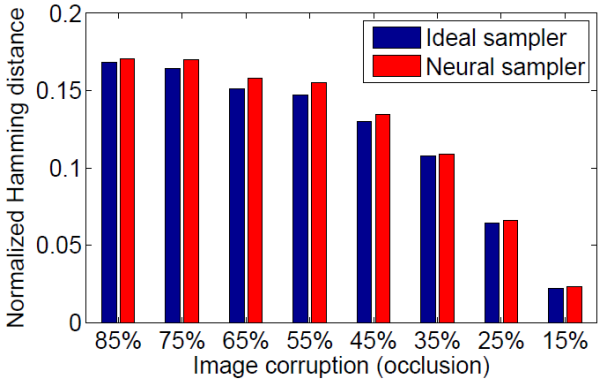

Fig. 19: Sampler generative performance analysis in terms of incorrectly generated pixels (Hamming distance).

Lastly, the mean HD for the TrueNorth RBM can be verified throughout the reconstruction process. In Fig. 20, convergence to the mean HD value for 35% image occlusion (dashed black line) occurs after about 10 reconstruction samples. This result is important to define the practical time expenditure demanded for the generative task of MNIST image reconstruction.

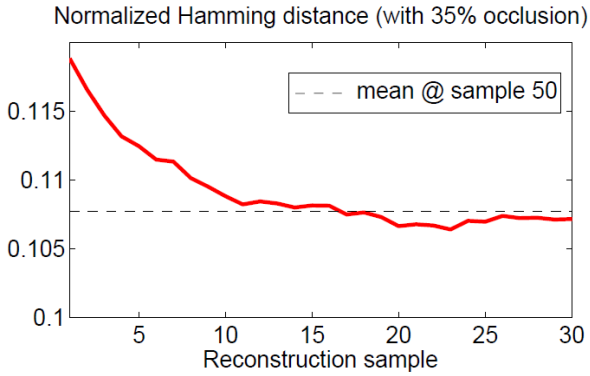

Fig. 20: Normalized Hamming distance during reconstruction of 35% occluded image on the TrueNorth RBM.

X. CONCLUSIONS AND FUTURE WORK

In this work, we have shown the first generative RBM implementation on neuromorphic hardware. For this, we followed a step-by-step procedure for producing the Gibbs sampling kernel -- the sigmoidal spiking probability -- using digital spiking neurons and for mapping the generative RBM algorithm onto a digital neuromorphic VLSI substrate. The neural sampler is an elegant solution as it uses bio-inspired dynamics to simultaneously incorporate the logistic function look-up and the comparison with a randomly generated number, which together represent a Gibbs sample. A discrete-time Markov chain (DTMC) analysis of the neural sampler was performed, resulting in a simplified method of obtaining the spiking probability without the need for long neuron behavior simulations. The generative performance of the neuromorphic adaptations were then verified using the Kullback-Leibler (KL) divergence and the Annealed Importance Sampling (AIS) algorithm. We also showed how mean squared error (MSE), along with the DTMC, can be used as an efficient method for obtaining insight into the sampler quality.

In the TrueNorth system, we followed a systematic development and implementation process of a modular architecture, which can be used for realizing generative and discriminative RBMs and DBNs on a substrate of digital neurosynaptic cores. The 3-stage architecture and the design automation procedure provide a path towards automated neural network applications on brain-inspired processors for more complex inference tasks, such as natural image recognition and time series generation. The modular characteristic of the architecture naturally lends itself to implementations of deeper networks (DBNs). Also, the architecture of stages 1 and 2 with the associated design automation procedure can even be used to realize other neural networks which are defined by sparse weight matrices. We are currently working on new algorithms which incorporate more of the hardware constraints during the training phase.

The developed architecture avails many of the features present in neuromorphic systems. The spike flow of the 3-stage architecture developed for the TrueNorth RBM uses spikes for communication between cores, propagating information between RBM layers. The number of computations is also reduced in the neuromorphic scenario as only non-zero multiplications are performed (i.e. only when a spike occurs does data processing take place), which is contrary to what occurs traditionally for matrix multiplications in CPUs. Additionally, the sampler makes use of stochastic neural properties to produce an approximate sigmoidal firing probability, necessary for the RBM sampling procedure. Despite these positive features, information processing in the network is somewhat sequential (i.e. basically two stages are being used at each instant), which is mainly a result of the limited weight values per neuron present in the hardware. Therefore, we are currently developing paths towards more continuous flows of spike processing in networks of neurosynaptic cores. A more continuous spiking neural sampling scheme would also be possible, however a slightly different neuron model (for example, with voltage-dependent leak and Gaussian noise [24]) and a hierarchical (tree) routing architecture -- as opposed to TrueNorth's mesh routing architecture with distributed cores -- would make an implementation more viable, as neural sampling would be continuous and constraints in network connectivity and weight values could possibly be relaxed.

ACKNOWLEDGEMENTS

The authors would like to thank all the members of the Brain-Inspired Computing Group at the IBM Almaden Research Center for their dedication and collaboration.

REFERENCES

[1] H. Larochelle and Y. Bengio, “Classification using discriminative restricted boltzmann machines,” in Proceedings of the 25th international conference on Machine learning, pp. 536–543, ACM, 2008.

[2] G. E. Dahl, D. Yu, L. Deng, and A. Acero, “Context-dependent pretrained deep neural networks for large-vocabulary speech recognition,” Audio, Speech, and Language Processing, IEEE Transactions on, vol. 20, no. 1, pp. 30–42, 2012.

[3] G. W. Taylor, G. E. Hinton, and S. T. Roweis, “Modeling human motion using binary latent variables,” Advances in neural information processing systems, vol. 19, p. 1345, 2007.

[4] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” Science, vol. 313, no. 5786, pp. 504–507, 2006.

[5] S. Haykin, Neural Networks and Learning Machines (3rd Edition). Prentice Hall, 2008.

[6] G. Hinton, S. Osindero, and Y.-W. Teh, “A fast learning algorithm for deep belief nets,” Neural computation, vol. 18, no. 7, pp. 1527–1554, 2006.

[7] G. Indiveri, B. Linares-Barranco, T. J. Hamilton, A. Van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Hafliger, S. Renaud, et al., “Neuromorphic silicon neuron circuits,” Frontiers in neuroscience, vol. 5, 2011.

[8] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, et al., “A million spiking-neuron integrated circuit with a scalable communication network and interface,” Science, vol. 345, no. 6197, pp. 668–673, 2014.

[9] M. Tommiska, “Efficient digital implementation of the sigmoid function for reprogrammable logic,” in Computers and Digital Techniques, IEEE Proceedings-, vol. 150, pp. 403–411, IET, 2003.

[10] A. Tisan, S. Oniga, D. MIC, and A. Buchman, “Digital implementation of the sigmoid function for fpga circuits,” ACTA TECHNICA NAPOCENSIS Electronics and Telecommunications, vol. 50, no. 2, p. 6, 2009.

[11] K. Lakshmi and M. Subadra, “A survey on fpga based mlp realization for on-chip learning,” Int. J. Sci. Eng. Res, pp. 1–9, 2013.

[12] S. Das, B. Pedroni, P. Merolla, J. Arthur, A. Cassidy, D. Modha, G. Cauwenberghs, and K. Kreutz-Delgado, “Gibbs Sampling with Low-Power Spiking Digital Neurons,” in IEEE Int. Symp. Circuits and Systems, 2015. Manuscript submitted for publication.

[13] J. G. Kemeny and J. L. Snell, Finite markov chains, vol. 356. Van Nostrand Princeton, NJ, 1960.

[14] A. Clark and S. Gilmore, “Terminating passage-time calculations on uniformised markov chains,” in Proceedings of the Twenty-Fourth annual UK Performance Engineering Workshop, pp. 64–75, Citeseer, 2008.

[15] A. S. Cassidy, P. Merolla, J. V. Arthur, S. K. Esser, B. Jackson, R. Alvarez-icaza, P. Datta, J. Sawada, T. M. Wong, V. Feldman,
A. Amir, D. B. dayan Rubin, E. Mcquinn, W. P. Risk, and D. S. Modha, “Cognitive computing building block: A versatile and efficient digital neuron model for neurosynaptic cores,” in International Joint Conference on Neural Networks (IJCNN). IEEE, 2013.

[16] Y. Tang and C. Eliasmith, “Deep networks for robust visual recognition,” in Proceedings of the 27th International Conference on Machine Learning (ICML-10), pp. 1055–1062, 2010.

[17] T. M. Cover and J. A. Thomas, Elements of information theory. John Wiley & Sons, 2012.

[18] R. M. Neal, “Annealed importance sampling,” Statistics and Computing, vol. 11, no. 2, pp. 125–139, 2001.

[19] R. Salakhutdinov and I. Murray, “On the quantitative analysis of deep belief networks,” in Proceedings of the 25th international conference on Machine learning, pp. 872–879, ACM, 2008.

[20] A. Amir, P. Datta, W. P. Risk, A. S. Cassidy, J. A. Kusnitz, S. K. Esser, A. Andreopoulos, T. M. Wong, M. Flickner, R. Alvarez-Icaza,et al., “Cognitive computing programming paradigm: a corelet language for composing networks of neurosynaptic cores,” in Neural Networks (IJCNN), The 2013 International Joint Conference on, pp. 1–10, IEEE, 2013.

[21] P. Merolla, J. Arthur, F. Akopyan, N. Imam, R. Manohar, and D. S. Modha, “A digital neurosynaptic core using embedded crossbar memory with 45pj per spike in 45nm,” in Custom Integrated Circuits Conference (CICC), 2011 IEEE, pp. 1–4, IEEE, 2011.

[22] S. K. Esser, A. Andreopoulos, R. Appuswamy, P. Datta, D. Barch, A. Amir, J. Arthur, A. Cassidy, M. Flickner, P. Merolla, et al., “Cognitive computing systems: Algorithms and applications for networks of neurosynaptic cores,” in Neural Networks (IJCNN), The 2013 International Joint Conference on, pp. 1–10, IEEE, 2013.

[23] T. Tieleman, “Training restricted boltzmann machines using approximations to the likelihood gradient,” in Proceedings of the 25th international conference on Machine learning, pp. 1064–1071, ACM, 2008.

[24] E. Neftci, S. Das, B. Pedroni, K. Kreutz-Delgado, and G. Cauwenberghs, “Event-driven contrastive divergence for spiking neuromorphic systems,” Frontiers in Neuroscience, vol. 7, p. 272, 2013.