

Temporal-Coded Deep Spiking Neural Network with Easy Training and Robust Performance

Shibo Zhou,¹ Xiaohua Li,¹ Ying Chen,² Sanjeev T. Chandrasekaran,³ Arindam Sanyal³

¹ Dept. of ECE, Binghamton University, Binghamton, NY, USA

² (Corresponding Author) Harbin Institute of Technology, China

³ Dept. of EE, University at Buffalo, Buffalo, NY, USA

{szhou19, xli}@binghamton.edu, yingchen@hit.edu.cn, {stannirk, arindams}@buffalo.edu

Abstract

Spiking neural network (SNN) is promising but the development has fallen far behind conventional deep neural networks (DNNs) because of difficult training. To resolve the training problem, we analyze the closed-form input-output response of spiking neurons and use the response expression to build abstract SNN models for training. This avoids calculating membrane potential during training and makes the direct training of SNN as efficient as DNN. We show that the nonleaky integrate-and-fire neuron with single-spike temporal-coding is the best choice for direct-train deep SNNs. We develop an energy-efficient phase-domain signal processing circuit for the neuron and propose a direct-train deep SNN framework. Thanks to easy training, we train deep SNNs under weight quantizations to study their robustness over low-cost neuromorphic hardware. Experiments show that our direct-train deep SNNs have the highest CIFAR-10 classification accuracy among SNNs, achieve ImageNet classification accuracy within 1% of the DNN of equivalent architecture, and are robust to weight quantization and noise perturbation.

1 Introduction

Spiking neural network (SNN) is interesting theoretically due to its strong bio-plausibility and practically because of its outstanding energy efficiency. Neurons communicate via spikes just as biological neurons. They work asynchronously, i.e., generate output spikes without waiting for all input neurons to spike. This leads to advantages such as spike sparsity, low latency, and high energy efficiency that are attractive for practical applications (Pfeiffer and Pfeil 2018; Tavanaei et al. 2019).

The performance of SNNs has fallen far behind conventional deep neural networks (DNNs). One of the primary reasons is that SNNs are difficult to train. DNNs are formulated with the standard layer response $\mathbf{y} = f(\mathbf{x}\mathbf{W} + b)$ where gradient backpropagation can be efficiently conducted. In contrast, for SNNs we have to simulate the temporal-domain neuron membrane potentials with non-differentiable spikes. Gradient evaluation is both difficult and time-consuming. Direct training of SNN has so far been limited to shallow networks only. No one has trained directly the SNNs on large datasets such as ImageNet.

SNNs are targeting neuromorphic hardware implementations. Their robustness on low-cost neuromorphic hardware with limited memory, high noise, and large parameter drift is critical for them to be competitive to DNNs in practical applications. This special yet important problem has been largely open.

In this paper, we resolve the difficult training problem by developing a direct-train framework that does not need calculating neuron membrane potentials. We address the robustness problem by developing a neuron circuit to evaluate timing jitter and by training deep SNNs under weight quantizations. Our major contributions are listed as follows.

- Closed-form analytical input-output responses of spiking neurons are studied systematically, which shows that the nonleaky integrate-and-fire neuron with single-spike temporal-coding is the best choice for direct-train SNNs.
- Deep SNNs such as SpikingVGG16 and Spiking-GoogleNet are developed and trained over the CIFAR-10 and ImageNet datasets. New benchmark results are obtained. To the best of our knowledge, this is the first time that direct-train SNN is reported for ImageNet.
- A phase-domain signal processing neuron circuit is designed to show that our neuron is more energy-efficient than others and is robust to input timing jitter and weight quantization. Besides, our deep SNNs are trained under weight quantization and noise perturbation to demonstrate their robustness.

This paper is organized as follows. Related works are introduced in Section 2. Spiking neurons and deep SNNs are described in Section 3. Experiments are presented in Section 4. Conclusions are given in Section 5.

2 Related Works

SNN training methods can be categorized into three classes: unsupervised learning, supervised learning with indirect training, and supervised learning with direct training (Pfeiffer and Pfeil 2018). For unsupervised learning, spike timing-dependent plasticity (STDP) is well known (Caporale and Dan 2008; Diehl and Cook 2015; Kheradpisheh et al. 2018; Lee et al. 2018). Nevertheless, the dependency on the local neuronal activities without a global supervisor makes it have low performance in deep networks.

For the second class, the popular approach is to translate pre-trained DNNs to SNNs (Tavanaei et al. 2019). This can be conducted by mapping DNN neuron values to SNN neuron spike rates (Diehl et al. 2015; Rueckauer et al. 2017; Sengupta et al. 2019) or spike times (Zhang et al. 2019). Although the translation approach has had the best performance among SNNs so far, it unfortunately sacrifices important SNN advantages such as spike sparsity and asynchronous processing, which degrades energy efficiency. Besides, many DNN techniques such as max-pooling and tanh activation are hard or inefficient to copy to SNN.

Among the third class, SpikeProp (Bohte, Kok, and La Poutre 2002) minimized the loss between the true and desired spike times with gradient descent rule over soft nonlinearity models. Gardner et al. (Gardner, Sporea, and Grüning 2015) applied a probability neuron model to calculate gradients. Gradient backpropagation was applied in (Hunsberger and Eliasmith 2016; Lee, Delbruck, and Pfeiffer 2016; Jin, Zhang, and Li 2018; Wu et al. 2019b) based on spike rate coding where gradient backpropagation through both SNN layer responses and neuron membrane potential dynamics was needed, which made the algorithms extremely complex. Backpropagation through SNN layer responses only was applied in (Mostafa 2017). Zhou and Li (Zhou and Li 2020) pointed out that the overly-nonlinear neuron response led to low training performance in SNNs. Existing direct training approaches still fall short of efficiency to deal with large datasets such as ImageNet.

A list of neuromorphic hardware has been developed for SNN, such as IBM TrueNorth (Merolla et al. 2014), Intel Loihi (Davies et al. 2018), and BrainScaleS (Aamir et al. 2018b). For energy efficiency, Hunsberger and Eliasmith (Hunsberger and Eliasmith 2016) estimated that a synaptic operation consumed only 8% of the energy of a microprocessor floating-point operation. Cao et al. (Cao, Chen, and Khosla 2015) showed that SNN implemented in a neuromorphic circuit with 45 pJ per spike was 185 times more energy-efficient than the FPGA-based DNN implementation. With 26 pJ per spike, IBM TrueNorth consumed 1.76×10^5 times less energy than computer simulations over microprocessors, and 769 times less energy than microprocessor-based neuromorphic hardware. Zhou et al. (Zhou et al. 2020) developed an analog neuron circuit with 19 pJ per spike, with which SNN-based YOYOv2 consumed only 0.247 mJ for detecting objects in an image frame.

Neuromorphic hardware has the problem of noise, parameter drifting, as well as severe limitation on resources such as memory. Implementing SNNs in neuromorphic hardware without robustness optimization showed heavy performance degradation (Esser et al. 2016; Gölztz et al. 2019). Rathi et al. (Rathi, Panda, and Roy 2018) studied the pruning of unimportant weights and quantizing important weights to improve energy efficiency, but for shallow SNNs only.

3 Deep SNN with Easy and Direct Training

The major hurdle to training SNNs is that every neuron’s membrane potential has to be calculated in each training iteration. This is not a problem when implementing SNNs in neuromorphic hardware but is computationally prohibitive




Figure 1: Integrate-and-fire spiking neuron model.

for software implementation and gradient-based training. Non-differential spike waveform is the second hurdle. Our objective is to completely avoid the calculation of membrane potential and spike waveform during training. Instead, we train the SNNs based on an abstract *layer response model* of the neurons. As illustrated in Fig. 1, the left figure is the spiking neuron model for real hardware implementation and inference only, where the weights w_{ji} need to be learned. The right figure is its layer response model, where we use the closed-form input-output response $t_j = f(t_i, w_{ji})$ to train the weights w_{ji} . Since the weights of the two figures are identical, the weights trained in the right figure are directly used to implement the SNN in the left figure.

In this section, we first show that only a special spiking neuron is appropriate for such direct training. Then, we show that this spiking neuron can be practically realized in energy-efficient hardware. Finally, based on this neuron, we propose a direct-train framework for deep SNNs.

3.1 Layer Response Models of Spiking Neurons

For the neuron illustrated in Fig. 1 (left), we consider the integrated-and-fire model. Membrane potential $v_j(t)$ of the neuron j is modeled as

$$\frac{dv_j(t)}{dt} + bv_j(t) = \sum_i w_{ji} \sum_k g(t - t_{ik}), \quad (1)$$

where b is a positive constant representing the leaky rate of the membrane potential, w_{ji} is the weight of the synaptic connection from the input neuron i to the output neuron j , $g(t)$ is the synaptic current kernel function or spike waveform, and t_{ik} is the spiking time of the k th spike of the i th input (pre-synaptic) neuron. $b > 0$ means *leaky integrate-and-fire* (LIF) neuron, while $b = 0$ means *nonleaky integrate-and-fire* (IF) neuron. Once $v_j(t)$ reaches spiking threshold θ , the neuron generates an output (post-synaptic) spike and the membrane potential is reset.

Information can be encoded in spike rate r_j , spike time t_j , or other means. We consider the first two, which we call rate coding and temporal coding, respectively. Rate r_j is the average spike rate from $t = 0$ to $t = T$. For temporal coding, each neuron generates a single spike during the time period T . We denote the spike time as t_j and adopt the time-to-first-spike (TTFS) code (Gölztz et al. 2019).

We have conducted a thorough study of the solutions to (1) in order to find desirable layer response models. Details are in Technical Appendix A. We analyze our major observations only in this subsection.

Consider rate coding first. With impulse spike $g(t) = \delta(t)$, nonleaky IF neuron has closed-form layer response

$$r_j = \text{ReLU} \left(\sum_i r_i \frac{w_{ji}}{\theta} \right), \quad (2)$$

where $\text{ReLU}(x) = \max\{0, x\}$. See (18) for derivations. Similar expressions exist for Heaviside and exponentially-decaying spike waveforms (21)(23). Since these expressions are identical to DNN's layer response, we can directly train a network implemented in software based on (2) and apply the resulted weights w_{ji} to the real SNN implemented in neuromorphic hardware.

Note that (2) is also the theoretical basis for translating DNNs to SNNs (Cao, Chen, and Khosla 2015; Diehl et al. 2015). It is interesting to see that direct-train SNN and translate-SNN become similar based on (2). The only difference is that the latter trains weights w_{ji}/θ instead of w_{ji} and thus needs weight normalization. Unfortunately, (2) is an approximate model only. Modeling error accumulates to the detrimental level in deep SNNs (Rueckauer et al. 2017). Corrections are developed for translate-SNNs to mitigate the error to some extent. The corrections need to calculate membrane potential, which makes direct-train difficult again.

For rate-coded LIF neurons, layer responses become numerically unstable for training. As shown in (25), LIF neuron with the impulse spike waveform has layer response

$$r_j = \text{ReLU}(-b \log^{-1}(1 - \theta / \sum_i w_{ji} / (e^{b/r_i} - 1))). \quad (3)$$

Random weights w_{ji} often make log function undefined, which means training can not proceed. The same problem happens for other spike waveforms (28)(30).

Next, for temporal coding, similar numerical instability happens for LIF neurons with the exponentially-decaying spike waveform. As shown in (42) and (44), the layer responses are expressed in either Lambert W function or quadratic equation roots. Random weights often result in negative or complex values that prevent gradient updating.

Fortunately, temporal-coded nonleaky IF neurons have layer responses desirable for direct training. With the exponentially-decaying spike waveform, the layer response can be formulated as (Mostafa 2017)

$$e^{\frac{t_j}{\tau}} = \sum_{i \in \mathcal{C}} e^{\frac{t_i}{\tau}} \frac{w_{ji}}{\sum_{\ell \in \mathcal{C}} w_{j\ell} - \theta} \quad (4)$$

where the set $\mathcal{C} = \{\forall k : t_k < t_j\}$. See (35) for derivations. With the Heaviside spike waveform, both IF and LIF neurons have similar layer responses, see (33) and (38). There is no significant modeling error and the expressions have nice numerical stability.

Since the energy efficiency of SNNs depends on the number of spikes, we prefer single-spike neurons. Therefore, the best choice is to adopt the single-spike temporal-coded nonleaky IF neuron to implement SNNs and train them based on (4). Note that we prefer the exponentially-decaying spike waveform rather than the Heaviside spike waveform because the former can be truncated to short spike duration in practice to enhance energy efficiency.

3.2 Circuit of Temporal-Coded Spiking Neuron

Most existing neuron circuits are designed for rate coding or for generating special spike patterns (Aamir et al. 2018a). Few circuits are reported for single-spike neurons. Therefore, we present a new neuron circuit design to show that the single-spike temporal-coded neuron with exponentially-decaying spike waveform can be realized practically. This circuit will also help us study energy efficiency, timing jitter, and weight quantization robustness.

For neuron circuit designs, prior works have used analog voltage domain (VD) signal processing to realize LIF neurons using CMOS circuits (Indiveri 2003; Wijekoon and Dudek 2009; Wu et al. 2015). VD signal processing techniques require high-gain amplifiers which are energy inefficient to design in advanced CMOS nodes with low intrinsic gain. To reduce energy consumption, we introduce phase-domain (PD) signal processing into SNN design. PD signal processing converts VD signal excursions into phase domain for further processing, and has well documented advantages over VD processing such as higher energy efficiency, ability to better leverage CMOS technology scaling and simpler circuit design. While PD signal processing has been widely adopted in the circuit design community for high energy-efficiency data-converter design (Sanyal et al. 2014; Sanyal and Sun 2016; Jayaraj et al. 2019a,b), a key contribution of this work is the introduction of PD signal processing to SNN design field for the first time.

Fig. 2 shows the circuit schematic of the IF neuron designed with voltage-controlled ring oscillator (VCO) based integrators. As shown in Fig. 2(a), if a voltage input $V_{in}(t)$ is applied to a ring VCO, its instantaneous phase is given by $\Phi(t) = \int 2\pi k_{vco} V_{in}(t) dt$, where k_{vco} is VCO tuning gain. Thus, a VCO acts as a perfect PD integrator with infinite dc gain (Sanyal and Sun 2016; Taylor and Galton 2010; Perrott et al. 2008), and is perfectly positioned to realize non-leaky integration. While the VCO phase is integral of its input, the phase output cannot be directly extracted. Rather, the voltage output of one of the inverters is read out which acts as a pulse-width modulated signal (PWM) and encodes the phase information in the width of its pulses as shown in Fig. 2(a). The ring VCO has a highly digital architecture and can operate from very low supply voltages. Hence, VCO can be used as an integrator in advanced CMOS processes which allows area and power scaling unlike VD integrators.

Fig. 2(b) shows the schematic of the proposed phase-domain IF neuron designed using VCOs, in which VCO1 and VCO3 act as global reference sources driven by constant inputs and are shared with all the neurons. T_1 and T_2 represent 2 inputs to the neuron while the resistors R_{1j} and R_{2j} denote the weights of synaptic connections of the inputs to the neuron, respectively. When the inputs to VCO2 spike, the output phase of the VCO2 jumps abruptly and VCO2 starts running at a high frequency. Outputs of VCO1 and VCO2 drive two counters which increment at the rising edge of VCO outputs. The difference between the two counters is used to drive VCO2 through negative feedback. The negative feedback loop forces VCO2 to track VCO1, and the frequency of VCO2 starts decaying exponentially till it reaches the frequency of VCO1 as shown in Fig. 2(b).

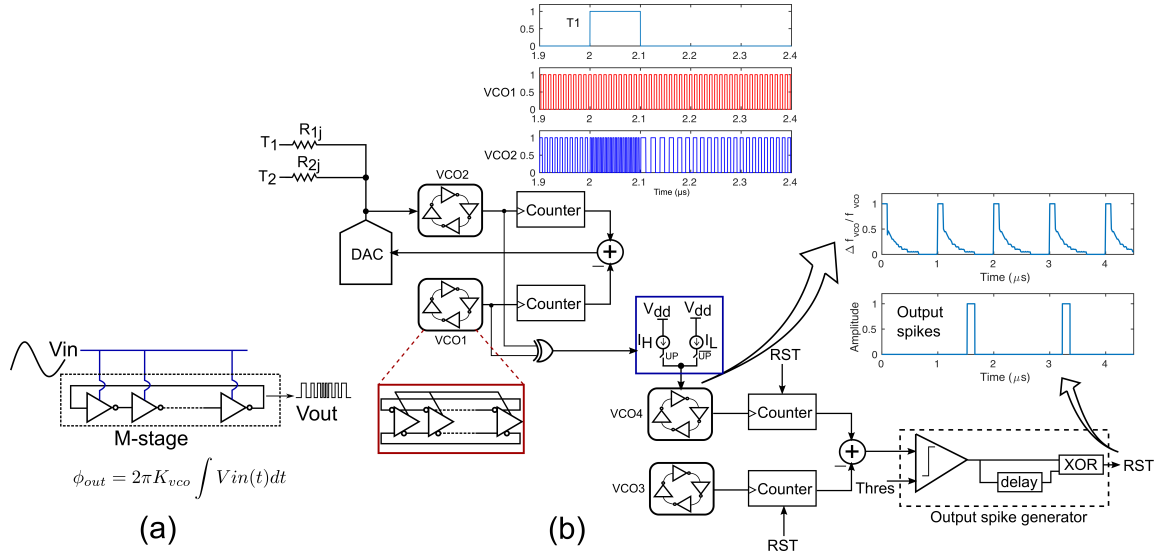


Figure 2: IF neuron circuit schematic using time-domain signal processing circuits. (a) Operation of ring-inverter-based VCO. (b) Circuit schematic of 1 neuron designed using ring VCOs. Waveforms of input, output and intermediate nodes show that T_1 generates an input spike every microsecond, VCO1/VCO2 reshape the spike into the exponentially-decaying spike waveform, and VCO3/VCO4 generate an output spike every two microseconds.

Thus, VCO2 realizes the exponential decaying synaptic kernel. The frequency difference between VCO1 and VCO2 is extracted using an XOR gate and sent to VCO4 for accumulation. VCO4 switches between two frequencies depending on whether the XOR output is high or low. Similar to VCO1 and VCO2, VCO3 and VCO4 also drive two counters at rising edges of their respective outputs. Once the difference between the two counter outputs exceeds a threshold value, an output spike is generated, and both the counters are reset.

The proposed circuit is highly scalable since it is built using digital CMOS circuits. The synaptic weights can be made tunable by using a digitally controllable resistor bank. The neuron consumes less than 10 pJ/spike in 65nm process and the energy consumption will reduce further with technology scaling. This leads to an energy efficiency gain of 45% over (Zhou et al. 2020).

3.3 Direct-Train Framework for Deep SNNs

For training, we implement the abstract SNNs in software based on (4). Specifically, in the ℓ th layer, $z_{\ell-1,i} = e^{t_{\ell-1,i}/\tau}$ and $z_{\ell,j} = e^{t_{\ell,j}/\tau}$ are used directly as neuron's input and output. For an L -layer deep SNN, define the input as \mathbf{z}_0 with elements $z_{0,i}$ and the final output as \mathbf{z}_L with elements $z_{L,i}$. Smaller $z_{L,i}$ means stronger classification output. Then we have $\mathbf{z}_L = f(\mathbf{z}_0; \mathbf{w})$ with nonlinear mapping f and trainable weight \mathbf{w} which includes all weights w_{ji}^ℓ . Let the targeting output be class c . We train the network with the loss function

$$\mathcal{L}(\mathbf{z}_L, c) = -\log \frac{z_{L,c}^{-1}}{\sum_{i \neq c} z_{L,i}^{-1}} + K \sum_{\ell=1}^L \sum_j \max \left\{ 0, \theta - \sum_i w_{ji}^\ell \right\} + \lambda \sum_{\ell=1}^L \sum_{j,i} (w_{ji}^\ell)^2. \quad (5)$$

The first term is to make $z_{L,c}$ the smallest (equivalently $t_{L,c}$ the smallest) one. The second term is the weight sum cost, which enlarges each neuron's input weight summation to increase its firing probability. The third term is L_2 regularization to prevent weights from becoming too large. The parameters K and λ are weighting coefficients. Thanks to the closed-form expression (4), gradient backpropagation can be used to train the weights. The training becomes nothing different from conventional DNNs.

Equations (4) and (5) were initially given in (Mostafa 2017). However, only shallow networks with 2 ~ 3 fully-connected layers were tried and the performance was low. One of the problems of the algorithm presented in (Mostafa 2017) is that $t_j > t_i$ for $i \in \mathcal{C}$ was not checked, which led to $t_j \leq t_i$ or even negative timing that was not hardware realizable. Another problem is that the algorithm is not suitable for deep SNN. As a result, this work did not arouse too much interest and the research progress has been slow.

We have resolved these problems by developing the following new algorithm to calculate (4): 1) Use the SORT function to sort inputs $e^{t_i/\tau}$; 2) Use the CUMSUM function to list all possible $\sum_{i \in \mathcal{C}} w_{ji}$ and $\sum_{i \in \mathcal{C}} w_{ji} e^{t_i/\tau}$; 3) Calculate $e^{t_j/\tau}$ for each element in the list; 4) Find the first one that satisfies $e^{t_j/\tau} > e^{t_i/\tau}$ for $i \in \mathcal{C}$ and $e^{t_j/\tau} \leq e^{t_i/\tau}$ for $i \notin \mathcal{C}$. Refer to our source code for details. This algorithm is much easier to implement and more efficient to run for deep SNNs. Note that real hardware SNN does not need this algorithm because spike time is naturally ordered.

As long as $\sum_i w_{ji} > \theta$, our algorithm can always provide hardware-compatible solutions $e^{t_j/\tau}$. With some heuristic convergence improving techniques (see experiment section), we have improved greatly both computation speed and convergence, which makes it possible to train deep SNNs.

To design deep SNNs, it is helpful to follow DNN architectures so as to exploit the extensive DNN research experience. For this we need to study how some DNN-specific techniques can be realized in SNNs, such as pooling, batch normalization (BN), and local response normalization (LRN). This issue has never been addressed in direct-train SNN including (Mostafa 2017).

For max-pooling, in software SNN training we have $e^{t_j/\tau} = \min_i e^{t_i/\tau}$. In real SNN hardware this can be realized as Fig. 1 using a large enough w_{ji} , which lets the first arrival spike activate the output spike. Average-pooling can be realized by letting $w_{ji} = \theta/(N - 1)$ for N inputs.

For BN, in software abstract SNN we need to calculate $e^{t_j/\tau} = \gamma/\sigma(e^{t_i/\tau} - \mu) + \beta$, where $\gamma, \sigma, \mu, \beta$ are parameters obtained from training. Rewriting it as

$$e^{t_j/\tau} = e^{(t_i + \tau \log \gamma/\sigma)/\tau} + e^{\log(\beta - \gamma\mu/\sigma)}, \quad (6)$$

it is easy to see that (6) can be realized in SNN hardware with two input neurons: one has spike time t_i with a constant delay $\tau \log \gamma/\sigma$, and the other has a constant spike time $\tau \log(\beta - \gamma\mu/\sigma)$. The weight should be $w_{ji} = \theta$. Note that spike time delay is a standard feature in SNN hardware such as IBM TrueNorth.

For LRN, the implementation is similar to BN but is much more complex. We also need to exploit the asynchronous property of SNN, i.e., reference timing $t = 0$ of a layer can be shifted to some non-zero value. In software SNN we have $e^{t_j/\tau} = e^{t_i/\tau}/(\gamma + \alpha \sum_{\ell} e^{2t_{\ell}/\tau})^{\beta}$ for constants γ, α and β . In hardware, we can use a neuron to calculate $\sum_{\ell} e^{2t_{\ell}/\tau}$ as average-pooling and then use another set of neurons to realize the rest operations similar to BN.

The above DNN-adaptation techniques would make our direct-train SNN look similar to translate-SNN. This is not surprising as our abstract SNN is designed to be trained similarly as DNN. The major differences are, first, we directly train the real SNN weights w_{ji} and thus do not need complex weight normalization or SNN structure change. Second, we have the flexibility to improve DNN techniques and train them to fit with SNN. For example, in BN, we can train SNN to guarantee the output in (6) be non-negative. Third, we can develop SNN specific techniques to outperform DNNs, such as adjusting θ in (5) to promote sparsity and energy efficiency. Note that besides spike sparsity, our temporal-coded SNN has neuron sparsity, i.e., only neurons in \mathcal{C} need to spike. Translate-SNN does not have such sparsity. The sparsity of translate-SNN usually means to reduce the number of spikes by using a shorter T , which makes the SNN work in transient response mode and thus suffer from performance degradation.

4 Experiments

In this section, we report our experiments over three standard datasets: MNIST, CIFAR-10, and ImageNet¹. We designed three deep SNN models, which are shown in Table 1, and evaluated their accuracy and robustness.

¹Our source code can be found at <https://github.com/zbs881314/Temporal-Coded-Deep-SNN>

4.1 Classification Accuracy

MNIST: The MNIST image pixels were normalized to $p_i \in [0, 1]$ and encoded into spiking time $t_i = \alpha(-p_i + 1)$. The parameter α was used to adjust spike temporal separation. We trained the network for 50 epochs with batch size 10 using the Adam optimizer. The learning rate started at 0.001 and gradually reduced to 0.0001 at the last epoch, with learning decay $\text{lr_decay} = (\text{learning_end} - \text{learning_start})/50$. We set $K = 100$ and $\lambda = 0.001$. According to (4), gradients could become very large in case $\sum_{\ell} w_{j\ell}$ is near θ , which is harmful to training. Therefore, we limited the maximum allowed row-normalized Frobenius norm of the gradient of each weight matrix to 10.

We trained this network with noisy input spike times. Classification accuracy is shown in Table 2. The proposed network had the highest accuracy (99.33%) among the SNNs listed in the table, yet had the smallest network size with the least number of trainable weights (21K). Because of the asynchronous operation of our SNN, on average 94% neurons spiked, which led to sparsity 0.94. The total consumed energy of the proposed network was 205 nJ assuming each spike cost 10 pJ based on the proposed neuron circuit and sparsity 0.94.

CIFAR-10: Our SNN model for CIFAR-10 was developed based on the VGG16 model (Liu and Deng 2015; Simonyan and Zisserman 2014), hence called SpikingVGG16. To encode image pixels to spike time, we applied encoding rule $t_i = \alpha p_i$ and used a relatively large α to enlarge spike time separation so that all pixel values could potentially be used. We exploited data augmentation such as crop, flip, and whiten to increase the diversity of data available for training. The learning rate started at 0.01 and ended at 0.0001. We ran 320 epochs and after 240 epochs the training tended to converge. The batch size was 128. The other hyper-parameters were the same as the MNIST experiments.

As shown in Table 2, the testing accuracy of our SpikingVGG16 set a new record at 92.68%, higher than all the listed SNNs including the state-of-the-art of (Wu et al. 2019b). Especially, our model had higher accuracy than the CNN-based VGG16 (Liu and Deng 2015). Our network enjoyed a sparsity of 0.62 which means only 62% neurons sent spikes. It consumed 0.336 mJ of energy for each image inference.

ImageNet: We built our deep SNN model based on the popular GoogleNet architecture (Szegedy et al. 2015), hence called SpikingGoogleNet. We simply replaced CNN layers with SCNN layers and FC layers with spiking FC layers. We used 1.2 million images to train the network. Training parameters were similar to the CIFAR-10 experiments, and the training procedure followed that of GoogleNet. The input image was $224 \times 224 \times 3$ and was randomly cropped from a resized image using the scale and aspect ratio augmentation. We used SGD with a batch size of 256, learning rate decay 0.0001, and momentum 0.9. We started from a learning rate of 0.1 and divided it by 10 three times. When the

MNIST	SCNN(5,32,2) → SCNN(5,16,2) → FC(10)
CIFAR-10	SpikingVGG16: SCNN(3,64,1) → SCNN(3,64,1) → MP(2) → SCNN(3,128,1) → SCNN(3,128,1) → MP(2) → SCNN(3,256,1) → SCNN(3,256,1) → SCNN(3,256,1) → MP(2) → SCNN(3,512,1) → SCNN(3,512,1) → SCNN(3,512,1) → MP(2) → SCNN(3,1024,1) → SCNN(3,1024,1) → SCNN(3,1024,1) → MP(2) → FC(4096) → FC(4096) → FC(512) → FC(10)
ImageNet	SpikingGoogleNet: replace GoogleNet CNN/FC layers with SCNN/FC layers

Table 1: Proposed models. SCNN(5,32,2) means spiking CNN layer with 32 5×5 kernels and stride 2. FC(10) means fully-connected layer with 10 output neurons. MP(2) means 2×2 max-pooling layer. $\tau = \theta = 1$. Bias is added as $t_0 = 0$.

Dataset	Models	Method	Accuracy	Weights(million)	Sparsity
MNIST	(Ciresan et al. 2011)	CNN	99.73%	0.069	no
	(Kheradpisheh et al. 2018)	SNN+STDP	98.40%	0.076	no
	(Lee et al. 2018)	SNN+STDP	91.10%	0.025	no
	(Mostafa 2017)	SNN+DT	97.55%	0.635	0.51
	(Zhang et al. 2019)	SNN+tran	99.08%	3.9	no
	(Wu et al. 2019a)	SNN+DT	99.26%	0.051	no
	Our Model	SNN+DT	99.33%	0.021	0.94
CIFAR-10	(Liu and Deng 2015) (VGG16)	CNN	91.55%	15	no
	(Hunsberger and Eliasmith 2016)	SNN+tran	82.95%	39	no
	(Rueckauer et al. 2017)	SNN+tran	90.85%	62	no
	(Sengupta et al. 2019) (VGG16)	SNN+tran	91.55%	15+	no
	(Wu et al. 2019b)	SNN+DT	90.53%	45	no
	Our Model	SNN+DT	92.68%	54	0.62
ImageNet	(Simonyan and Zisserman 2014) (VGG16)	CNN	71.5%	138	no
	(Szegedy et al. 2015) (GoogleNet)	CNN	69.8%	6.8	no
	(Rueckauer et al. 2017) (VGG16)	SNN+tran	49.61%	138	no
	(Sengupta et al. 2019) (VGG16)	SNN+tran	69.96%	138+	no
	(Zhang et al. 2019) (VGG16)	SNN+tran	70.87%	138+	no
	Our Model	SNN+DT	68.8%	6.8	0.56

Table 2: Classification accuracy comparison. DT: direct training. tran: translate SNN.

training error was no longer reducing, retraining and fine-tuning with very small learning rates were conducted until the test accuracy no longer increased.

Results of Top-1 testing accuracy are shown in Table 2. Our SpikingGoogleNet achieved 68.8% accuracy, only 1.0% lower than the CNN-based GoogleNet. VGG16-based SNNs had slightly better performance because CNN-based VGG16 had better accuracy than GoogleNet. But they had much larger network sizes.

For energy consumption, our network enjoyed a sparsity of 0.56 which means only 56% of neurons were activated during image inference. The total energy consumption was thus 0.038 mJ. For the translate-SNN model (Sengupta et al. 2019), if implemented in IBM TrueNorth with 26 pJ per spike and having an average of 256 spikes per neuron with rate coding, each of them would consume 900 mJ, several orders bigger than ours.

To the best of our knowledge, few works were reported over the CIFAR-10 dataset using direct training SNN (except (Wu et al. 2019b)) and none was reported over the ImageNet dataset. Our models hence set up new benchmark accuracy in both cases. There is still a big gap between SNN and DNN (Tan and Le 2019). Nevertheless, our experiments showed that our proposed SNNs could achieve similar ac-

curacy as the DNNs with similar network size and architecture. This fact was also demonstrated in many translate SNN works. We expect that SNN performance could catch up rapidly after the training hurdle is resolved.

4.2 Robustness

Neuron Robustness: For neuromorphic circuits, a problem is that they are very sensitive to changes in operating conditions such as supply voltage or temperature. In Fig. 2, we have used differential architecture in which VCO phase is always compared with phase from a reference VCO rather than using the absolute phase from a single VCO. While this architectural choice doubles area and energy consumption, it also increases robustness to changes in operating conditions.

For temporal-coded neurons, an important distortion measure is spike timing jitter, which is defined as $|t_j^{\text{measured}} - t_j^{\text{desired}}|/t_j^{\text{desired}}$. For mixed digital-analog circuits, there are quantization noise and circuit noise, which we simply combine into synaptic weight quantization noise. We conducted circuit simulation under various input timing jitters and weight quantizations, and the results are summarized in Table 3. Both input jitter and quantization of synaptic weights change the temporal location of the output spikes. As the input jitter was swept from 1% to 9%, the output jitter var-

Input Jitter (%)	Output Jitter (%)	Quantization Level	Output Jitter (%)
1	0.024	4-bit	0.957
3	0.189	8-bit	0.268
5	0.397	16-bit	0.008
7	0.789	24-bit	0.004
9	0.997	32-bit	0.004

Table 3: Jitter in output spike versus input jitter and synaptic weight quantization.

ied from 0.024% to 0.997% only. Such a good inherent jitter suppression capability was due to the phase-locked loop (PLL) which low-pass-filtered input jitter.

Quantizing synaptic weights (for 2-input case) increased output jitter from 0.004% in the case of 32-bit quantization to 0.957% in the case of 4-bit quantization. Coarse quantization thus led to large jitter. One of the reasons was that in our neuron model (4) the denominator $\sum_{\ell} w_{j\ell} - \theta$ could be small and a slight change of weights $w_{j\ell}$ might cause a big change in t_j . To mitigate this problem, we can set θ in (5) to a bigger number during training.

Examining the numbers in Table 3 more carefully, we see that 4-bit quantization, which has quantization SNR (signal to noise ratio) $6.02 \times 4 \approx 24$ dB, led to output jitter of 0.957% which means SNR $20 \log_{10}(1/.00957) \approx 40$ dB. A 40 dB SNR means that the neuron was extremely robust to weight quantization.

Deep SNN Robustness: To evaluate the robustness of deep SNNs, we experimented with both weight quantization and noise perturbation. For weight quantization, the weights were quantized to 32-bit, 8-bit, 4-bit, and 2-bit words. Thanks to our easy training models, we could retrain the deep SNNs simply following the procedure developed for conventional CNN quantization (Li, Zhang, and Liu 2016; Rastegari et al. 2016). Specifically, the forward inference used quantized weights while the backward gradient propagation used full-precision weights. We first trained with 32-bit quantization. After the training converged, we applied 8-bit quantization and retrain the SNNs. This procedure was repeated until the 2-bit quantization.

Table 4 shows the testing accuracy under weight quantization. For the models over MNIST, CIFAR10 and ImageNet, weight quantization caused the worst accuracy loss of 0.22%, 1.75%, and 8.8%, respectively. As comparison, we listed two typical CNN weight quantization results (Cheng et al. 2018; Zhang et al. 2018), which indicated a similar performance degradation pace. The results demonstrated that the SNN models were robust to weight quantization for relatively small datasets and small networks. For larger networks, the weights would better be encoded in 4-bit or over.

To evaluate noise perturbation, we added random noise to the trained weights. Experiment results are shown in Fig. 3, together with the results of weight quantization in Table 4 (expressed in SNR). We find that 24dB quantization noise (4-bit quantization) reduced ImageNet classification accuracy to 65.2%. Noise at 24dB SNR reduced ImageNet clas-

Model	Bits	MNIST	CIFAR-10	ImageNet
Our Models	32	99.33%	92.68%	68.8%
	8	99.32%	91.87%	66.1%
	4	99.21%	91.38%	65.2%
	2	99.11%	90.93%	60.0%
(Cheng et al. 2018)	32	98.66%	84.80%	-
	8	98.48%	84.07%	-
	2	96.34%	81.56%	-
(Zhang et al. 2018)	32	-	92.1%	70.3%
	4	-	-	70.0%
	2	-	91.8%	68.0%

Table 4: Accuracy versus weight quantization.

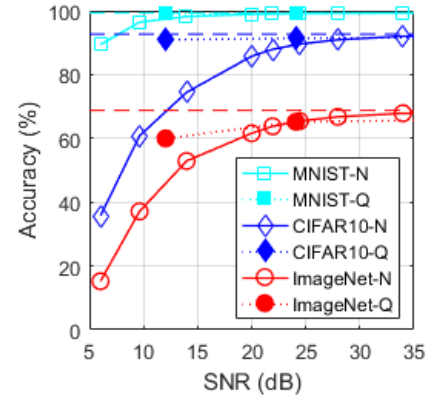


Figure 3: Accuracy versus weight quantization (Q) and noise perturbation (N). Dashed Lines: baseline (32-bit quantization and noiseless cases). Note: 8,4,2-bit quantizations correspond to quantization SNRs 48, 24, 12dB, respectively.

sification accuracy to 65.43%. Both cases had a small 3.5% performance loss only.

Because Table 3 indicated that neuron noise (explained as jitter) was usually very small, we just need to pay attention to high SNR scenarios, such as 24dB and above (corresponding to 4-bit or above quantization). In this case, the accuracy reduction was negligible. Therefore, the deep SNNs were robust to both weight quantization and noise perturbation.

Extra experiment results are in Technical Appendix B.

5 Conclusion

In this paper, we develop direct-train deep SNNs that can be easily trained over large datasets such as ImageNet. We show that our SNNs can realize classification accuracy within 1% of (or even better than) DNNs of similar size and architecture. A circuit schematic of the adopted neuron is designed with phase-domain signal processing, which shows 45% energy efficiency gain over the existing state of the art. Both the neuron and the deep SNNs are demonstrated as robust to timing jitter, weight quantization, and noise. The easy training, high performance, and robustness indicate that SNNs can be competitive to DNNs in practical applications.

References

- Aamir, S. A.; Müller, P.; Kiene, G.; Kriener, L.; Stradmann, Y.; Grübl, A.; Schemmel, J.; and Meier, K. 2018a. A mixed-signal structured adex neuron for accelerated neuromorphic cores. *IEEE transactions on biomedical circuits and systems* 12(5): 1027–1037.
- Aamir, S. A.; Stradmann, Y.; Müller, P.; Pehle, C.; Hartel, A.; Grübl, A.; Schemmel, J.; and Meier, K. 2018b. An accelerated lif neuronal network array for a large-scale mixed-signal neuromorphic architecture. *IEEE Transactions on Circuits and Systems I: Regular Papers* 65(12): 4299–4312.
- Bohte, S. M.; Kok, J. N.; and La Poutre, H. 2002. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing* 48(1-4): 17–37.
- Cao, Y.; Chen, Y.; and Khosla, D. 2015. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision* 113(1): 54–66.
- Caporale, N.; and Dan, Y. 2008. Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31: 25–46.
- Cheng, H.-P.; Huang, Y.; Guo, X.; Huang, Y.; Yan, F.; Li, H.; and Chen, Y. 2018. Differentiable fine-grained quantization for deep neural network compression. *arXiv preprint arXiv:1810.10351* .
- Ciresan, D. C.; Meier, U.; Gambardella, L. M.; and Schmidhuber, J. 2011. Convolutional neural network committees for handwritten character classification. In *2011 International Conference on Document Analysis and Recognition*, 1135–1139. IEEE.
- Davies, M.; Srinivasa, N.; Lin, T.-H.; China, G.; Cao, Y.; Choday, S. H.; Dimou, G.; Joshi, P.; Imam, N.; Jain, S.; et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38(1): 82–99.
- Diehl, P. U.; and Cook, M. 2015. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in computational neuroscience* 9: 99.
- Diehl, P. U.; Neil, D.; Binas, J.; Cook, M.; Liu, S.-C.; and Pfeiffer, M. 2015. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *2015 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Esser, S.; Merolla, P.; Arthur, J.; Cassidy, A.; Appuswamy, R.; Andreopoulos, A.; Berg, D.; McKinstry, J.; Melano, T.; Barch, D.; et al. 2016. Convolutional Networks for Fast, Energy-Efficient Neuromorphic Computing. *CoRR abs/1603.08270*. <http://arxiv.org/abs/1603.08270> .
- Gardner, B.; Sporea, I.; and Grüning, A. 2015. Learning spatiotemporally encoded pattern transformations in structured spiking neural networks. *Neural computation* 27(12): 2548–2586.
- Göltz, J.; Baumbach, A.; Billaudelle, S.; Breitwieser, O.; Dold, D.; Kriener, L.; Kungl, A. F.; Senn, W.; Schemmel, J.; Meier, K.; et al. 2019. Fast and deep neuromorphic learning with time-to-first-spike coding. *arXiv preprint arXiv:1912.11443* .
- Hunsberger, E.; and Eliasmith, C. 2016. Training spiking deep networks for neuromorphic hardware. *arXiv preprint arXiv:1611.05141* .
- Indiveri, G. 2003. A low-power adaptive integrate-and-fire neuron circuit. In *IEEE International Symposium on Circuits and Systems*, volume 4, IV–IV.
- Jayaraj, A.; Danesh, M.; Chandrasekaran, S. T.; and Sanyal, A. 2019a. Highly Digital Second-Order $\Delta\Sigma$ VCO ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers* 66(7): 2415–2425.
- Jayaraj, A.; Das, A.; Arcot, S.; and Sanyal, A. 2019b. 8.6fJ/step VCO-Based CT 2nd-Order $\Delta\Sigma$ ADC. In *IEEE Asian Solid-State Circuits Conference (A-SSCC)*, 197–200.
- Jin, Y.; Zhang, W.; and Li, P. 2018. Hybrid macro/micro level backpropagation for training deep spiking neural networks. In *Advances in neural information processing systems*, 7005–7015.
- Kheradpisheh, S. R.; Ganjtabesh, M.; Thorpe, S. J.; and Masquelier, T. 2018. STDP-based spiking deep convolutional neural networks for object recognition. *Neural Networks* 99: 56–67.
- Lee, C.; Srinivasan, G.; Panda, P.; and Roy, K. 2018. Deep spiking convolutional neural network trained with unsupervised spike-timing-dependent plasticity. *IEEE Transactions on Cognitive and Developmental Systems* 11(3): 384–394.
- Lee, J. H.; Delbruck, T.; and Pfeiffer, M. 2016. Training deep spiking neural networks using backpropagation. *Frontiers in neuroscience* 10: 508.
- Li, F.; Zhang, B.; and Liu, B. 2016. Ternary weight networks. *arXiv preprint arXiv:1605.04711* .
- Liu, S.; and Deng, W. 2015. Very deep convolutional neural network based image classification using small training sample size. In *2015 3rd IAPR Asian conference on pattern recognition (ACPR)*, 730–734. IEEE.
- Merolla, P. A.; Arthur, J. V.; Alvarez-Icaza, R.; Cassidy, A. S.; Sawada, J.; Akopyan, F.; Jackson, B. L.; Imam, N.; Guo, C.; Nakamura, Y.; et al. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345(6197): 668–673.
- Mostafa, H. 2017. Supervised learning based on temporal coding in spiking neural networks. *IEEE transactions on neural networks and learning systems* 29(7): 3227–3235.
- Perrott, M.; et al. 2008. A 12-bit 10-MHz bandwidth continuous-time ADC with a 5-bit 950-MS/s VCO-based quantizer. *IEEE J. Solid-State Circuits* .
- Pfeiffer, M.; and Pfeil, T. 2018. Deep learning with spiking neurons: opportunities and challenges. *Frontiers in neuroscience* 12: 774.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, 525–542. Springer.

- Rathi, N.; Panda, P.; and Roy, K. 2018. STDP-based pruning of connections and weight quantization in spiking neural networks for energy-efficient recognition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 38(4): 668–677.
- Rueckauer, B.; Lungu, I.-A.; Hu, Y.; Pfeiffer, M.; and Liu, S.-C. 2017. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in neuroscience* 11: 682.
- Sanyal, A.; Ragab, K.; Chen, L.; Viswanathan, T.; Yan, S.; and Sun, N. 2014. A hybrid SAR-VCO $\Delta\Sigma$ ADC with first-order noise shaping. In *IEEE Custom Integrated Circuits Conference*, 1–4.
- Sanyal, A.; and Sun, N. 2016. A 18.5-fJ/step VCO-based 0–1 MASH $\Delta - \Sigma$ ADC with digital background calibration. In *2016 IEEE Symposium on VLSI Circuits (VLSI-Circuits)*, 1–2. IEEE.
- Sengupta, A.; Ye, Y.; Wang, R.; Liu, C.; and Roy, K. 2019. Going deeper in spiking neural networks: Vgg and residual architectures. *Frontiers in neuroscience* 13: 95.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* .
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- Tan, M.; and Le, Q. V. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946* .
- Tavanaei, A.; Ghodrati, M.; Kheradpisheh, S. R.; Masquelier, T.; and Maida, A. 2019. Deep learning in spiking neural networks. *Neural Networks* 111: 47–63.
- Taylor, G.; and Galton, I. 2010. A mostly-digital variable-rate continuous-time delta-sigma modulator ADC. *IEEE Journal of Solid-State Circuits* 45(12): 2634–2646.
- Wijekoon, J. H.; and Dudek, P. 2009. A CMOS circuit implementation of a spiking neuron with bursting and adaptation on a biological timescale. In *IEEE Biomedical Circuits and Systems Conference*, 193–196.
- Wu, J.; Chua, Y.; Zhang, M.; Yang, Q.; Li, G.; and Li, H. 2019a. Deep spiking neural network with spike count based learning rule. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–6. IEEE.
- Wu, X.; Saxena, V.; Zhu, K.; and Balagopal, S. 2015. A CMOS spiking neuron for brain-inspired neural networks with resistive synapses and in-situ learning. *IEEE Transactions on Circuits and Systems II: Express Briefs* 62(11): 1088–1092.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; Xie, Y.; and Shi, L. 2019b. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1311–1318.
- Zhang, D.; Yang, J.; Ye, D.; and Hua, G. 2018. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, 365–382.
- Zhang, L.; Zhou, S.; Zhi, T.; Du, Z.; and Chen, Y. 2019. Tdsnn: From deep neural networks to deep spike neural networks with temporal-coding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1319–1326.
- Zhou, S.; Chen, Y.; Li, X.; and Sanyal, A. 2020. Deep SCNN-based Real-time Object Detection for Self-driving Vehicles Using LiDAR Temporal Data. *IEEE Access* .
- Zhou, S.; and Li, X. 2020. Spiking Neural Networks with Single-Spike Temporal-Coded Neurons for Network Intrusion Detection. *arXiv preprint arXiv:2010.07803* .