

Exploring Extreme Quantization in Spiking Language Models

Malyaban Bal, Yi Jiang, Abhranil Sengupta
School of Electrical Engineering and Computer Science
The Pennsylvania State University
University Park, PA 16802, USA
Email: {mjb7906, yijiang, sengupta}@psu.edu

Abstract—Despite the growing prevalence of large language model (LLM) architectures, a crucial concern persists regarding their energy and power consumption, which still lags far behind the remarkable energy efficiency of the human brain. Recent strides in spiking language models (LM) and transformer architectures aim to address this concern by harnessing the spiking activity of biological neurons to enhance energy/power efficiency. Doubling down on the principles of model quantization and energy efficiency, this paper proposes the development of a novel binary/ternary (1/1.58-bit) spiking LM architecture. Achieving scalability comparable to a deep spiking LM architecture is facilitated by an efficient knowledge distillation technique, wherein knowledge from a non-spiking full-precision “teacher” model is transferred to an extremely weight quantized spiking “student” LM. Our proposed model represents a significant advancement as the first-of-its-kind 1/1.58-bit spiking LM, and its performance is rigorously evaluated on multiple text classification tasks of the GLUE benchmark.

Index Terms—Neuromorphic Computing, Language Modelling, Spiking Neural Networks (SNNs), Model Quantization, Knowledge Distillation.

I. INTRODUCTION

Large language models (LLMs) are becoming increasingly popular due to their wide-ranging applications in various natural language processing (NLP) tasks. The immense scale of transformer [1] based language models (LMs) comes at the cost of increased energy/power consumption. By integrating bio-plausible models with extreme parameter quantization techniques, we can markedly diminish energy and power costs, thereby fostering a sustainable future for deep neural models.

Spiking neural networks (SNNs) [2] emulate biological processes, where communication between neurons takes the form of spikes. This sparse, event-driven flow of information enables efficient computation and communication in neuromorphic hardware, leading to significant energy savings [3]. Though substantial work on SNN-based transformer architectures are done on vision-based datasets, recent works have explored spiking LM architectures for language generation [4] and classification tasks [5].

While spiking architectures assist in reducing multi-bit neuron activations to binary spiking activity, they offer limited support in minimizing parameter overhead. Model quantization techniques stands out as a promising solution, effectively slashing the memory footprint and computational expenses of expansive models, all while upholding their competitive

performance. Model quantization can occur during training or post-training. While post-training quantization [6] is simpler, it often leads to more significant accuracy degradation compared to techniques that incorporate quantization-aware training. Our work is primarily geared towards developing a quantization-aware training approach.

The spiking architectures’ capacity to distribute computations across the temporal domain is pivotal in overcoming the formidable challenge of crafting an exceptionally quantized spiking LM. By amalgamating scalable spiking LM with extreme model quantization [7], [8], this endeavor strives to markedly diminish model size and enhance energy/power efficiency through the utilization of spiking activation and 1/1.58-bit weights. This quantized spiking architecture is adaptable for implementation on neuromorphic chips and specialized hardware accelerators such as “In-Memory” binary neural network (BNN) accelerators [9]. Moreover, executing the binarized spiking LM over an optimal number of time-steps enables us to reach near full-precision accuracies while capitalizing on the specialized “In-Memory” hardware accelerator tailored for BNNs [9].

In our paper, we introduce a framework for training a spiking LM with parameters quantized to 1/1.58-bits. We leverage the technique of implicit differentiation at equilibrium for training the spiking architecture, as outlined in previous works [5], [10]. This approach offers outstanding memory efficiency during training compared to Backpropagation Through Time (BPTT), which demands substantial memory for storing a large computational graph. Moreover, it eliminates the need for surrogate gradient methods by implicitly computing gradients, thus addressing the non-differentiability challenge inherent in training spiking models with BPTT. Within specific constraints [11], this learning paradigm parallels biologically plausible and energy-efficient training methods, such as equilibrium propagation [12]–[14], thereby reinforcing a neuromorphic perspective on learning.

In this paper, leveraging the average spiking rate at equilibrium of the quantized spiking LM, we use an efficient knowledge distillation (KD) technique to transfer knowledge from a non-spiking high-precision “teacher” architecture to a spiking 1/1.58-bit “student” LM. This KD methodology is pivotal for the creation of an extremely quantized spiking LM, enabling efficient training of our “student” model even with

limited resources. For this study, we adopt the encoder-based BERT [15] architecture as our LM, chosen for its suitability in text classification tasks due to its ability to capture bi-directional contextual information.

The primary contributions of this paper are as follows:

- We propose an efficient framework for training-aware extreme model quantization in spiking LM architectures.
- We leverage the equilibrium dynamics of the spiking LM to efficiently perform the model compression and quantization through the use of KD.
- To the best of our knowledge, this is the first-of-its-kind 1/1.58-bit spiking LM which is evaluated on multiple NLP tasks of the GLUE benchmark.

II. METHODS

In this section, we begin by exploring the architecture of our 1/1.58-bit BERT-based spiking LM. Subsequently, we provide a concise overview of the training mechanism and the effective KD technique employed, enabling us to achieve such extreme quantization by distilling knowledge from a full-precision “teacher” model.

A. Architecture and Learning Dynamics

The base architecture of this model follows previous BERT-based spiking architectures [5], and comprises of stacked spiking encoder layers. Each spiking encoder layer comprises of a spiking attention module and multiple intermediate layers as described in Fig 1. In this highly quantized version, we substitute all linear projection layers, typically full precision feed-forward layers, with quantized binary/ternary linear layers, whose operation is detailed later. Additionally, given the spiking nature of this architecture, communication between neuronal layers exclusively takes place in the form of spikes (Fig. 1). The dynamics of the spiking neurons are given as:

$$\begin{aligned} u_i[t + \delta] &= \gamma u_i[t] + \sum_j (w_{ij} s_j[t]) + b_i, \\ s_i[t + 1] &= S(u_i[t + \delta]), \\ u_i[t + 1] &= u_i[t + \delta] - V_{th} s_i[t + 1] \end{aligned} \quad (1)$$

Here, γ denotes the leaky term; V_{th} is the spiking threshold; $u_i[t]$ and $s_i[t]$ signifies the membrane potential and spike from the i^{th} neuron at time t ; w_{ij} represents the synaptic weight between the pre and post-synaptic neurons; $t + \delta$ denotes an intermediate time step to determine if the neuron has fired; b_i indicates a bias term; S denotes the non-differentiable function for spike generation, with subtraction as reset operation.

The average spiking rate (ASR) for each layer i , which can be defined as a weighted-average function as follows, $a_i[t] = \frac{\sum_{\tau=1}^t \gamma^{t-\tau} s_i[\tau]}{\sum_{\tau=1}^t \gamma^{t-\tau}}$, is leveraged during training and during KD following [5] and is given as,

$$a_i[t + 1] = \frac{1}{V_{th}} (W_i(a_{(i-1)}[t + 1]) + b_i - \frac{u_i[t + 1]}{\sum_{i=0}^t \gamma^i}) \quad (2)$$

For linear layers, Eqn. 2 can be derived from Eqn. 1. Theoretically, as $t \rightarrow \infty$ the ASRs eventually converge to equilibrium

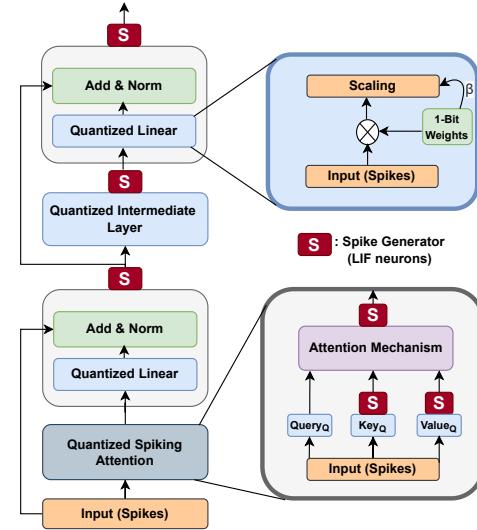


Fig. 1. High-level architecture of each encoder layer of the 1-bit Spiking transformer architecture with Quantized Linear layers. All linear layers, including those used in the attention module, are quantized to 1-bit.

and we can derive steady-state equations for ASR of linear layers and formulate surrogate steady-state functions for non-linear layers [5], [10] of the form:

$$a_i^* = \sigma\left(\frac{1}{V_{th}}(W_i(a_{i-1}^*) + b_i)\right) \quad (3)$$

where clipping function $\sigma(x)$ bounds the values within [0,1] and W_i is operation of the corresponding layer. The layer-wise convergence dynamics of ASRs of the 1-bit quantized Spiking LM model is shown in Fig. 2a. We leverage only the ASR values at equilibrium to compute error-gradients using implicit differentiation at equilibrium [10] as shown below,

$$\frac{\partial L(a^*)}{\partial \theta} = -\frac{\partial L(a^*)}{\partial a^*} (J_{g_\theta}^{-1}|_{a^*}) \frac{\partial f_\theta(a^*)}{\partial \theta} \quad (4)$$

where, L is the loss function used, $g_\theta(a) = f_\theta(a) - a$, f is the steady-state equation of ASR, J^{-1} is the inverse Jacobian of g_θ when $a = a^*$, i.e., at equilibrium. Hence, unlike BPTT, we eliminate the need to store all the intermediate states. Additionally, we circumvent the need for surrogate gradients to address the non-differentiability issue encountered by BPTT during SNN training, as we employ ASRs (real values) rather than spikes during training. This capacity to utilize ASRs during training is also crucial for formulating an extremely quantized spiking architecture, as detailed later.

B. Quantizing linear layers to 1/1.58-bit

The predominant energy expenditure during inference of a spiking LM arises from the numerous linear projection layers utilized across various components, including the attention module and intermediate layers such as feed-forward and output layers. This is primarily attributed to the floating-point matrix multiplication operation performed by these layers, which incurs significantly higher energy costs compared to normalization or residual layers [7]. Linear projection operations

involve two tensors: neuronal activity and synaptic weights. Primarily, previous research has typically concentrated on extreme quantization (< 2 bits) of either neuronal activity [5] or synaptic weights [7], [16]. In contrast, our objective is to simultaneously quantize both using just 1 bit (neurons are intrinsically quantized due to spiking behavior). We further explore ternary weight quantization, utilizing a $\log_2(3)$ i.e. 1.58-bit representation, where the states correspond to -1/0/1, in order to bridge the accuracy gap with the full-precision model even further.

In line with [7], [8], for 1-bit quantization, the weights (W) are first centred around zero-mean and then the model weights are binarized to either +1 or -1 using the signum function. Thus, after binarization the quantized weight is,

$$W_{Q_{1-bit}} = \text{Sign}(W - \alpha) \quad (5)$$

where, $\text{Sign}(W_{ij}) = \begin{cases} +1, & \text{if } W_{ij} > 0, \\ -1, & \text{if } W_{ij} \leq 0. \end{cases}$ and $\alpha = \frac{1}{nm} \sum_{ij} W_{ij}$ is the mean of the weight matrix $W \in \mathbb{R}^{n \times m}$. For the 1.58-bit implementation, the quantization can be done as follows:

$$W_{Q_{1.58-bit}} = \text{Round}(\text{Clip}(W/(\beta + \epsilon), -1, 1)) \quad (6)$$

where, ϵ is a small constant. The output of the linear layer is scaled by a constant $\beta = \frac{1}{nm} \sum_{ij} |W_{ij}|$. In addition to the extreme quantization, our spiking LM distinguishes itself from non-spiking designs by quantizing the input to linear layers, solely by binary spikes, contrasted with continuous values [7], [16], [17]. It is crucial to emphasize that we harness this spiking activity during inference. During training, as detailed in the previous section, we utilize the full-precision ASR which helps us maintain the theoretical constraints of quantization [7]. Even with extreme quantization, comparable performance to full-precision model is attributed to the fact that the spiking LM operates over T timesteps, effectively spreading neuron activation precision across the temporal dimension. Additionally, as illustrated in Fig. 2b, the ASR convergence dynamics of the quantized model aligns closely with that of the full-precision model, indicating no additional overhead on the operating latency of the spiking LM.

C. Knowledge Distillation (KD)

Training large spiking transformer models, such as those found in LMs, demand significant resources. Our approach, which depends on training-aware extreme quantization, exacerbates the challenge of training from scratch. Building on prior research [5] regarding KD from ANN to SNN models, our goal is to streamline the training process. We achieve this by employing a full-precision “teacher” LM to impart knowledge to a quantized spiking “student” LM. This approach enables efficient compression of knowledge and significantly reduces the required resources compared to training from scratch. This formulation, akin to Eqn. 4, capitalizes on the steady-state ASR of the output of intermediate layers of the quantized spiking LM. It utilizes this information to minimize

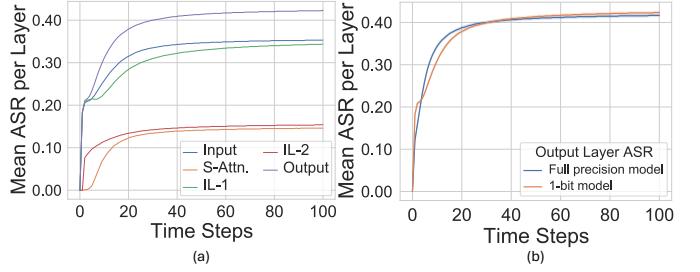


Fig. 2. Results obtained after passing a set of randomly sampled inputs from MRPC dataset. (a) The convergence dynamics of the different sub-layers of an encoder layer of the 1-bit SpikingBERT4 model. (b) Comparison of output layer ASR convergence dynamics of full precision (weights) and 1-bit SpikingBERT4. Y-axis in both shows mean (over number of neurons) of the ASR while X-axis shows the operating time steps.

a loss function, incorporating the full precision activity of corresponding intermediate layers within a “teacher”. This mechanism facilitates the effective transfer of knowledge leveraging the intermediate layers—specifically, the output of each individual encoder block—of both the “teacher” and the “student” model. The underlying loss is given as,

$$L_{h_i} = \text{MSE}(\text{ASR}(S_{h_i}^*)W_p, T_{f(h_i)}) \quad (7)$$

Here, MSE is mean squared error loss function, $\text{ASR}(S_{h_i}^*)$ represents the equilibrium ASR of the output neurons in the i^{th} SE layer of the “student” model, while $T_{f(h_i)}$ denotes the output of the $f(h_i)^{th}$ layer in the “teacher” model. W_p signifies a linear projection aligning the dimensionality of the “student” layer with its corresponding layer in the “teacher”. The function f maps the “student” layer h_i to a specific target layer in the “teacher” network.

III. RESULTS

In this section, we showcase the efficacy of extreme quantization techniques primarily on text classification tasks of the General Language Understanding Evaluation (GLUE) benchmark [18]. The experiments were conducted on Nvidia RTX A5000 GPUs (8 units), each equipped with 24GB of memory.

A. Datasets

We used Quora Question Pair (QQP), Microsoft Research Paraphrase Corpus (MRPC) to evaluate our model’s performance on similarity and paraphrase tasks. For inference-oriented evaluations, we opted for Multi-Genre Natural Language Inference (MNLI), Question-answering NLI (QNLI) datasets. For single-sentence sentiment analysis tasks, we utilized the Stanford Sentiment Treebank (SST-2). For all tasks, we keep the maximum sequence length at 128.

B. Experimental Setup

All experiments use four 1/1.58-bit spiking encoder blocks (Fig. 1) where encoding dimension of the tokens in the input is 768 and the intermediate size of the model is 3072. Training starts with a full-precision BERT-based spiking LM [5], pre-trained on the Wikipedia corpus. Subsequently, we conduct

TABLE I
RESULTS SHOWING PERFORMANCE (ACCURACY) OF OUR QUANTIZED SPIKING MODELS AGAINST SOME STANDARD MODELS AND OTHER EFFICIENT IMPLEMENTATIONS OF BERT ON GLUE EVALUATION SET.

Model	QQP	MNLI	SST-2	QNLI	MRPC
CBoW [18]	75.0	57.1	79.5	62.5	75.0
REM W2-A4 [16]	75.7	58.3	82.9	75.3	-
BinaryBERT _{50%} [19]	66.7	39.2	54.1	59.5	68.3
TernaryBERT [19]	74.1	32.7	53.1	59.3	68.3
BERT ₅ + PF [20]	84.1	67.7	81.6	80.9	78.6
NAS-BERT ₅ + KD [20]	85.8	74.4	87.3	84.9	79.6
BERT _{TINY} Adam [5]	81.1	65.3	80.1	77.8	69.9
1-bit SpikingBERT₄	83.8	75.4	86.7	80.5	75.8
1.58-bit SpikingBERT₄	85.4	77.1	87.1	83.1	77.3
SpikingBERT ₄	86.8	78.1	88.2	85.2	79.2

task-specific internal layer KD for each individual dataset. To quantize the weights, we employ Eqn. 5 for 1-bit and Eqn. 6 for 1.58-bit variant, followed by intermediate layer KD as per Eqn. 7 with non-spiking full-precision fine-tuned BERT as “teacher”. After KD, we refine our quantized models by minimizing cross-entropy loss against the true labels. This facilitates the training of our quantized models tailored to the specific task. Since this is the first work of exploring extreme quantization in spiking LM, our results are compared with some common NLP models and other efficient implementation of BERT. The results are demonstrated in Table. I. The model trained with ternary (1.58-bit) weights outperform that trained with binary (1-bit) weights and closely match the performance of full-precision SpikingBERT. Internal-layer KD (Eqn. 7) is crucial for high accuracy in quantized models. Without it, quantization-aware training on the proposed spiking LM leads to a significant 6-8% accuracy drop across the datasets.

C. Energy & Power Efficiency

As demonstrated empirically in Fig. 2b for the output layer, the spiking activity of neurons post-quantization remains comparable to that of full-precision models. Furthermore, empirical results suggest that the total number of $Norm\#OPS$ in full-precision (floating point ACC) and 1/1.58-bit models (Integer ACC) are similar; for example, for the MRPC dataset, the ratio of $Norm\#OPS$ of 1.58-bit SpikingBERT to full-precision SpikingBERT is 1.06 (i.e., ≈ 1). The total normalized OPS [9] is defined as $Norm\#OPS = \frac{\sum_i IFR_i * Layer\#OPS_{i+1}}{\sum Layer\#OPS}$, where IFR_i represents the number of spikes over inference time steps averaged across the number of neurons. However, it is worth noting here that each accumulative operation in 1/1.58-bit quantized models are at least an order of magnitude more energy efficient than full-precision models [9] (for instance, a preliminary estimate in terms of integer ACC versus floating point ACC energy consumption in 45nm CMOS technology yields $9\times$ energy efficiency [21]). Thus, the combination of extreme weight quantization and spiking neuronal activity enables a remarkable reduction in the model’s size, energy and power consumption.

IV. CONCLUSION & FUTURE WORKS

In our paper, we delve into leveraging the equilibrium-based convergence dynamics of spiking architectures to craft an extremely quantized spiking LM. This approach drastically shrinks model size, facilitating deployment on edge-based computing devices with limited resources. Furthermore, using computationally efficient “In-Memory” BNN accelerators over time enables us to seamlessly perform inference time tradeoff of accuracy and energy in complex sequence-based NLP tasks.

Future endeavors can broaden the application of this technique to encompass other deep spiking architectures, including GPT-like decoder-based LMs. Additional research efforts can aim to minimize the accuracy disparity between the full-precision models and quantized spiking transformer-based architectures. This technique of extreme quantization can also be explored in other deep spiking architectures beyond LMs.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the U.S. National Science Foundation under award No. CAREER #2337646, CCSS #2333881, CCF #1955815, and EFRI BRAID #2318101 and by Oracle Cloud credits and related resources provided by the Oracle for Research program.

REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” 2017.
- [2] S. Ghosh-Dastidar and H. Adeli, “Spiking neural networks,” *International journal of neural systems*, vol. 19, no. 04, pp. 295–308, 2009.
- [3] A. Sengupta, Y. Ye, R. Wang, C. Liu, and K. Roy, “Going deeper in spiking neural networks: VGG and residual architectures,” *Frontiers in neuroscience*, vol. 13, 2019.
- [4] R.-J. Zhu, Q. Zhao, G. Li, and J. K. Eshraghian, “SpikeGPT: Generative Pre-trained Language Model with Spiking Neural Networks,” 2023.
- [5] M. Bal and A. Sengupta, “Spikingbert: Distilling bert to train spiking language models using implicit differentiation,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 10, 2024, pp. 10998–11006.
- [6] J. Zhang, Y. Zhou, and R. Saab, “Post-training quantization for neural networks with provable guarantees,” *SIAM Journal on Mathematics of Data Science*, vol. 5, no. 2, pp. 373–399, 2023.
- [7] H. Wang, S. Ma, L. Dong, S. Huang, H. Wang, L. Ma, F. Yang, R. Wang, Y. Wu, and F. Wei, “Bitnet: Scaling 1-bit transformers for large language models,” *arXiv preprint arXiv:2310.11453*, 2023.
- [8] Z. Liu, B. Oguz, A. Pappu, L. Xiao, S. Yih, M. Li, R. Krishnamoorthi, and Y. Mehdad, “Bit: Robustly binarized multi-distilled transformer,” *Advances in neural information processing systems*, vol. 35, pp. 14 303–14 316, 2022.
- [9] S. Lu and A. Sengupta, “Exploring the connection between binary and spiking neural networks,” *Frontiers in neuroscience*, vol. 14, 2020.
- [10] M. Xiao, Q. Meng, Z. Zhang, Y. Wang, and Z. Lin, “Training feedback spiking neural networks by implicit differentiation on the equilibrium state,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 14 516–14 528, 2021.
- [11] S. Bai, J. Z. Kolter, and V. Koltun, “Deep equilibrium models,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [12] B. Scellier and Y. Bengio, “Equilibrium propagation: Bridging the gap between energy-based models and backpropagation,” *Frontiers in computational neuroscience*, vol. 11, p. 24, 2017.
- [13] M. Bal and A. Sengupta, “Sequence learning using equilibrium propagation,” *arXiv preprint arXiv:2209.09626*, 2022.
- [14] J. Lin, M. Bal, and A. Sengupta, “Scaling snns trained using equilibrium propagation to convolutional architectures,” *arXiv preprint arXiv:2405.02546*, 2024.

- [15] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [16] H. Bai, L. Hou, L. Shang, X. Jiang, I. King, and M. R. Lyu, “Towards efficient post-training quantization of pre-trained language models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 1405–1418, 2022.
- [17] S. Ma, H. Wang, L. Ma, L. Wang, W. Wang, S. Huang, L. Dong, R. Wang, J. Xue, and F. Wei, “The era of 1-bit llms: All large language models are in 1.58 bits,” *arXiv preprint arXiv:2402.17764*, 2024.
- [18] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, “GLUE: A multi-task benchmark and analysis platform for natural language understanding,” *arXiv preprint arXiv:1804.07461*, 2018.
- [19] H. Qin, Y. Ding, M. Zhang, Q. Yan, A. Liu, Q. Dang, Z. Liu, and X. Liu, “Bibert: Accurate fully binarized bert,” *arXiv preprint arXiv:2203.06390*, 2022.
- [20] J. Xu, X. Tan, R. Luo, K. Song, J. Li, T. Qin, and T.-Y. Liu, “Nas-bert: task-agnostic and adaptive-size bert compression with neural architecture search,” in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 1933–1943.
- [21] S. Han, J. Pool, J. Tran, and W. Dally, “Learning both weights and connections for efficient neural network,” in *Advances in neural information processing systems*, 2015, pp. 1135–1143.