SpikeSim Evaluation

Spiking Neural Network Hardware Performance Analysis

Abhineet Agarwal (22B1219

IIT Bombay
EE746 Neuromorphic Engineering

November 28, 2025

Outline

SpikeSim Repository Overview

Repository Structure

- SNN training & quantization framework
- Hardware evaluation tools (ELA)
- NICE hardware inference module

Key Features

- VGG9 architecture for SNNs
- Direct encoding (non-spiking input)
- 4-bit weight quantization support

GitHub: Intelligent-Computing-Lab/SpikeSim

Implementation Workflow

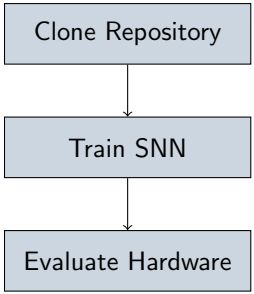

Task 1: FTJ vs RRAM Device Comparison

Device Parameters

- **FTJ:** $R_{on} = 1\text{M}\Omega$, $R_{off} = 10\text{M}\Omega$ (10:1 ratio) — Ferroelectric Tunnel Junction
- **RRAM:** $R_{on} = 20\text{k}\Omega$, $R_{off} = 200\text{k}\Omega$ (10:1 ratio) — Resistive Random Access Memory

Performance Comparison Results

Metric	FTJ	RRAM	Difference
Area (μm^2)	8,422,356	8,425,529	-0.04%
Energy (pJ)	15,140,653	15,906,934	-4.82%
Latency (ns)	62,832,640	62,832,640	0.00%

Both devices: 1-bit per cell — VGG9 architecture — 64×64 crossbar

Task 2: CIFAR-10 Training Results

Training Configuration

- Architecture: VGG9 with direct encoding
- Time steps: $T = 5$
- Quantization: 4-bit weights
- Learning rate: 10^{-3}
- Batch size: 128

Accuracy Progression

Epoch	Accuracy (%)
2	47.96
6	65.97
10	76.65
14	79.89
18	80.29
20	**80.30**

Final Test Accuracy

80.30%

Steady improvement with convergence at epoch 20

20 epochs — SGD optimizer — Leaky membrane (decay=0.5)

Task 2: CIFAR-100 Training Results

Training Configuration

- Architecture: VGG9 with direct encoding
- Time steps: $T = 10$ (increased)
- Quantization: 4-bit weights
- Learning rate: 10^{-3}
- Batch size: 128

Training Dynamics

Epoch	Accuracy (%)
2	10.63
6	26.23
10	34.83
14	41.76
16	**42.28**
20	42.28

Final Test Accuracy

42.28%

Convergence achieved at epoch 16

100 classes — Convergence at epoch 16 — Higher complexity vs CIFAR-10

NICE Hardware Inference

Inference Workflow

- Load pre-trained 4-bit quantized VGG9 model
- Configure hardware parameters (ADC precision, timesteps)
- Evaluate hardware-aware inference accuracy

Implementation In Progress

BinOp class dependency issue identified
Requires hw_models_new.py module integration

NICE: Neuromorphic In-memory Computing Engine

Summary & Key Findings

Device Comparison

- FTJ shows 4.8% energy reduction
- Minimal area difference (0.04%)
- Identical latency performance

Training Results

- CIFAR-10: 80.30% accuracy
- CIFAR-100: 42.28% accuracy
- 4-bit quantization effective

Implementation

- Created `ela_spikesim_ftj.py` with FTJ device model
- Extended ELA with energy and area calculations
- Configured training scripts for CIFAR-10/100
- Implemented 4-bit quantization with SGD optimizer

SpikeSim provides comprehensive SNN evaluation framework for hardware-aware design

Thank You!

Questions?