Graduate School of Science and Technology Master's Thesis Abstract

Laboratory name (Supervisor)	Dependable System (Michiko Inoue (Professor))		
Student ID	2411420	Submission date	2026 / 7 / 15
Name	RAHMAN MIZANUR		
Thesis title	**Effective Quantized Neural Networks on Faulty Memristor Crossbar Arrays**		

Abstract

Memristor crossbar arrays offer a compelling path toward energy-efficient neural network acceleration, yet permanent stuck-at faults (SAFs), conductance quantization errors, and stochastic device noise collectively undermine inference reliability. Existing mitigation strategies typically rely on expensive per-chip retraining or redundant hardware overhead, leaving the concurrent effects of SAFs, conductance quantization, and write noise largely unaddressed. To overcome these challenges, we propose **FILA-QAT (Fault-Injected Level-Adjusted Quantization-Aware Training)**, a unified framework that enforces fault constraints directly within the discrete integer-level domain to achieve inherently fault-tolerant, quantization-efficient, and noise-resilient inference. Additionally, a three-sigma-standard-deviation weight-clamping rule is established during the fault-injection warm-up phase to stabilize accuracy across different network architectures. During deployment, the framework seamlessly applies a **level adjustment (LA)** mechanism within the mapping phase. This ensures that noise robustness is fully preserved when programmed conductances are perturbed by log-normal write noise, completely bypassing the need for per-device retraining or hardware redundancy. Evaluated across multiple network architectures under fault rates ranging from 5% to 30%, FILA-QAT consistently delivers superior performance, yielding up to a 3.21% accuracy gain over post-training quantization baselines while maintaining high absolute inference accuracy of up to 97.93% on MLP/MNIST, 99.65% on LeNet-5/MNIST, and 87.79% on AlexNet/CIFAR-10 under severe 30% fault conditions. Furthermore, when deployed via our **weight-bipartite matching with level adjustment (WBM+LA)** pipeline, the framework yields up to a 6.33% accuracy recovery over naive sequential mapping. Finally, under stochastic log-normal write noise (σ = 0.05), the proposed framework completely neutralizes baseline degradation, reducing the maximum accuracy drop from 0.18% down to 0.00%.