

Enabling Fast Deep Learning on Tiny Energy-Harvesting IoT Devices

Sahidul Islam*, Jieren Deng[†], Shanglin Zhou[†], Chen Pan[‡], Caiwen Ding[†], Mimi Xie*

*Department of Computer Science, The University of Texas at San Antonio

[†]Department of Computer Science, University of Connecticut

[‡]Department of Computer Science, Texas A&M University–Corpus Christi

Abstract—Energy harvesting (EH) IoT devices that operate intermittently without batteries, coupled with advances in deep neural networks (DNNs), have opened up new opportunities for enabling sustainable smart applications. Nevertheless, implementing those computation and memory-intensive intelligent algorithms on EH devices is extremely difficult due to the challenges of limited resources and intermittent power supply that causes frequent failures. To address those challenges, this paper proposes a methodology that enables fast deep learning with low-energy accelerators for tiny energy harvesting devices. We first propose *RAD*, a resource-aware structured DNN training framework, which employs block circulant matrix and structured pruning to achieve high compression for leveraging the advantage of various vector operation accelerators. A DNN implementation method, *ACE*, is then proposed that employs low-energy accelerators to profit maximum performance with small energy consumption. Finally, we further design *FLEX*, the system support for intermittent computation in energy harvesting situations. Experimental results from three different DNN models demonstrate that *RAD*, *ACE*, and *FLEX* can enable fast and correct inference on energy harvesting devices with up to 4.26X runtime reduction, up to 7.7X energy reduction with higher accuracy over the state-of-the-art.

Index Terms—Deep Learning, Low-energy Accelerator, Energy Harvesting, IoT

I. INTRODUCTION

In this IoT era where all things trend towards to be embedded with electronics, DNN brings new opportunities for embedded IoT devices to be smarter, more versatile, and able to handle more complex jobs. With DNN algorithms implemented on device, data can be processed locally instead of transmitting large amount of data to the cloud server for further processing. Such a procedure not only significantly reduces processing pressure on the cloud but also improves energy efficiency and time for decision making. Resource-constrained embedded IoT devices however face three major challenges from computation, memory, and energy perspectives for implementing DNNs. First, embedded IoT devices have significantly less computational units and lower CPU frequency (e.g. 1-16MB) compared with high-performance computers. Since DNNs are computationally expensive, executing DNNs on embedded IoT devices results in significantly long execution time. Second, embedded IoT devices are equipped with much smaller (e.g. 1-16Kb) memory due to the size constraint. Third, majority IoT devices are powered with batteries, which can be depleted in days or even hours, requiring high maintenance cost which is especially undesirable if devices are implanted or deployed in harsh environment.

Energy harvesting (EH), which can scavenge energy from ambient environment, has become a promising technology to provide energy supply for IoT devices. However, the power

supply can be intermittent and low. Since DNN execution is time and energy-consuming, the power outage is prone to happen frequently in the middle of the execution. As a result, implementing DNN in energy harvesting devices is a highly challenging job because the implementation requires consideration of both the resource constraints and the intermittent power supply which complicates the development of such systems. Recent work proposed software system SONIC and TAILS [7] to enable fast deep learning inference on battery-less systems. Those systems exploit the special structured and loop-heavy computations of DNN and enable their intermittent executions by continuously saving the loop control states to the non-volatile memory after each instruction. However, correctness of intermittent execution requires special mechanism to handle the inconsistency issues.

To our knowledge, none of the existing work has considered the various vector operations supported by the low-energy accelerators in the training process and thus fail to efficiently leverage those accelerators in implementation and make efficient progress with existing software support for intermittent execution. To generate the best framework for the DNN implementation, we need to design the DNN model based on the resource constraints and employ the compression methods that can take advantage of those hardware accelerators while considering their limitations and the constraints of device. Besides, to avoid the progress setback due to power failures while guaranteeing correctness, we need special support for those vector operations executed on accelerators. To this end, this paper proposes the first framework that enables fast deep learning leveraging low-energy accelerators on resource-constrained EH IoT devices with three contributions as follows:

- A resource-aware structured DNN training framework, *RAD*, is proposed that consists of four different components: architecture search, compression, normalization, and fixed point calculation.
- A software system, *ACE*, is designed that efficiently implements DNN algorithm, including acceleration-aware data flow design, hardware acceleration, overflow-aware computation and circular buffer convolution.
- A novel on-demand robust checkpointing scheme, *FLEX*, is proposed which can minimize or completely avoid progress setbacks, even if a power outage happens.

The remainder of this paper is organized as follows: Section II provides the background and related work; Section III provides the framework overview and proposes the techniques for deep learning on energy harvesting devices; Section IV presents the experimental evaluation, and Section V concludes this work.

II. BACKGROUND AND RELATED WORK

Energy Harvesting Resource-Constrained Devices: Typical microcontrollers of IoT edge devices have minimal memory size and compute resources. Even a small DNN model like LeNet that has millions of operations takes several minutes to execute [7]. As a result, implementing DNN algorithms is more challenging [23]. DNN algorithms need to be redesigned and optimized to fit those small devices with good accuracy and run time. Besides, the energy needs to be harvested to a certain threshold to begin computation. However, the energy drained fast, resulting in frequent power failures, which incurs a problem with two dimensions. First, resuming the computation without wasted work. Second, checkpointing must not consume much energy and time. As a result, there are also opportunities to re-think about the traditional checkpointing mechanism.

Low Energy Accelerators: To efficiently deploy DNN algorithms on energy harvesting devices, vector operation accelerators like TI's low-energy accelerator (LEA) [17] can perform vector operations such as FFT(Fast Fourier Transform), IFFT (Inverse FFT), MAC (Multiply and Accumulation), ADD, etc. Such low energy vector processor performs vector operations without any CPU intervention.

DNN Pruning: Compared with other pruning solutions of DNN algorithms, block-circulant matrix (BCM) based DNNs [5] and structured pruning [21] [22] serves as good candidates for energy harvesting resource-constrained devices.

a) *Structured Pruning:* Structured pruning is proposed to structurally remove entire filters, channels, or filter shapes from the weight matrix [6]. Structured pruning becomes more hardware-friendly by taking advantage of the pruned weight matrices with regular shapes [20], avoiding introducing extra indices that indicate the pruned locations.

b) *Block-circulant matrix:* Block-circulant matrix (BCM) based DNN algorithm can drastically reduce memory footprint and computation pressure, consequently achieving low energy consumption. Since its implementation requires Fourier and inverse Fourier transformation, we can employ the FFT/IFFT accelerators for implementation.

The key idea of block-circulant matrix (BCM) based fully connected (FC) layers is to partition the original weight matrix $\mathbf{W} \in \mathbb{R}^{m \times n}$ into blocks of square sub-matrices, and each sub-matrix is a circulant matrix. Specifically, the matrix-vector multiplication can be implemented via "FFT→element-wise multiplication→IFFT", i.e., $\text{IFFT}(\text{FFT}(\mathbf{p}_{ij}) \circ \text{FFT}(\mathbf{x}_j))$ using the BCM format. For the inference phase, the computational complexity of this FC layer is $O(pqk \log k)$. Similarly, the storage complexity is $O(pqk)$ because only \mathbf{w}_{ij} or $\text{FFT}(\mathbf{w}_{ij})$ of each sub-matrix needs to be stored, which is equivalent to $O(n)$ for small p, q values. Therefore, the simultaneous acceleration and model compression can be achieved.

Related Works: Several works have implemented CNN on IoT devices. SONIC is an intermittence-aware software system with specialized support for DNN inference [7]. NeuroZERO introduces a co-processor architecture consisting of the main microcontroller that executes scaled-down versions of a (DNN) inference task [11]. A software/hardware co-design technique

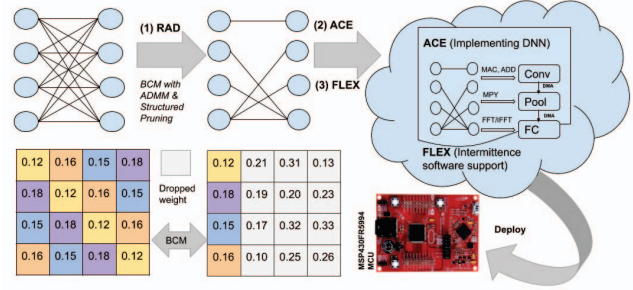


Fig. 1: System Overview.

that builds an energy-efficient low bit trainable system is proposed in [4]. Efficient memory aware data flow designs for inference are proposed [8]. Intermittent-Aware Neural Architecture Search for task based inference and Hardware Accelerated Intermittent inference is proposed in [13] and [9], respectively. ePerceptive is on-device CPU-based inference [14], based-on SONIC [7]. Different from the above works, this is the first work that explores BCM-based DNN algorithms on resource-constrained energy harvesting IoT devices.

III. DEEP LEARNING ON ENERGY HARVESTING DEVICE

System overview: We propose a resource-aware training, pruning, and implementation framework that takes various types of vector operations supported by low-energy accelerators into consideration. Figure 1 shows the system overview of the proposed framework which consists of three techniques named *RAD*, *ACE*, and *FLEX*, as well as the responsibility of each technique, and their correlation. Resource-aware DNN training and pruning method, *RAD*, provides the resource-aware model, compression, normalization, and fixed point calculation. *ACE* then implements DNN algorithm on EH device with efficient data flow design, circular buffer convolution, and hardware accelerators. A on-demand hybrid checkpointing method, *FLEX*, that also takes advantage of the special structure of DNN helps deal with the frequent power failures while avoiding progress setbacks.

A. Resource-aware DNN Pruning (RAD)

RAD aims to train and prune a DNN model that is aware of the available device resources. *RAD* starts with a backbone model with good accuracy by doing architecture search. Then *RAD* performs resource-aware compression with small accuracy loss while achieving less memory footprint and latency. *RAD* trains the model offline. Figure 2 shows the typical computing and memory resources in an IoT device. Besides, CPU, LEA can be used for executing vector operations of DNN algorithms, while SRAM and FRAM are used as data storage. SRAM is smaller with lower dynamic read-write energy compared with FRAM. SRAM is used as a buffer for intermediate data, where FRAM is used for storing the DNN model and the control data for intermittent computation, which will be discussed in Section III.C.

Modeling challenges: There are several challenges during DNN modeling. 1) SRAM/FRAM size 2) CPU frequency and inference time 3) Overflow error 4) Model accuracy. Aware of

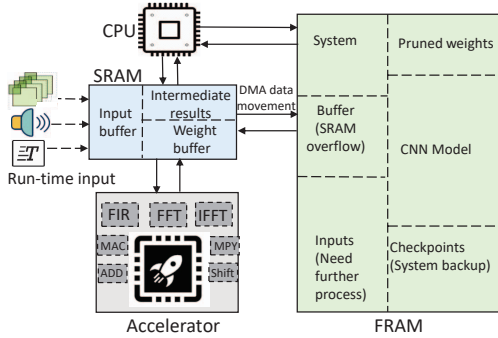


Fig. 2: Resource Overview in Embedded Devices.

the above challenges, *RAD* trained the DNN model offline. The model must fit into the FRAM with acceptable inference time and accuracy. *RAD*'s architecture search technology finds a suitable model and further compresses it. As we deal with fixed-point calculation in a resource constraint device, there is possible data overflow during vector operations like MAC and FFT. As a result, *RAD* performs several normalizations so that data stays between ranges during calculation.

Fixed-point quantization: Traditionally, DNN models are trained using high precision floating-point calculation. During on-device deployment, this is costly in terms of latency and memory. In this work, we use low-precision fixed-point representation. To implement the low-precision fixed-point representation, *RAD* maps high precision floating-point to 16-bits fixed-point so that *ACE* can deploy the DNN model with fixed-point calculation.

Normalization: For normalization, *RAD* first sets the data range G_i with a minimum value, G_{min} , as -1 and a maximum value, G_{max} , as 1. *RAD* then normalizes the data within this range of $[-1, 1]$. Finally, to avoid the value of the computed intermediates exceeding this range during inference, *RAD* uses the cosine normalization [12] to constrain the values of the computed intermediates into $[-1, 1]$.

TABLE I: BCM compression for 512*512 fully connected layer

Kernel Size	Block size	Compressed kernel	Storage reduction
1048576 Byte	16	65536 Byte	93.75%
	32	32768 Byte	96.87%
	64	16384 Byte	98.43%
	128	8192 Byte	99.21%
	256	4096 Byte	99.60%

Compression: The fully connected (FC) layers are partitioned into 2D blocks of square sub-matrices and each of sub-matrices is a circulant matrix. This block circulant matrix (BCM) [5] enables the FFT-based computation for the FC layers. In addition, the structured pruning is also applied on the convolutional (CONV) layers to reduce the redundancy of CONV. The combination of BCM on FC and structured pruning on CONV can achieve drastically reduced memory footprint with a very small accuracy drop. Especially for the BCM, shown in Table I, the memory footprint reduction for FC layers is up to 99.60%.

Alternating Direction Method of Multipliers (ADMM)-Regularized Optimization: The structured pruning problem can be effectively solved using the ADMM-regularized optimization. After the BCM-based pretrained DNN model is obtained, we define the constraint for the structured pruning separately on *accelerators* via the ADMM-regularized optimization [3].

Consider an N -layer DNN, where the weights and biases in the i -th CONV layer are \mathbf{W}_i and \mathbf{b}_i . The optimization problem is shown below.

$$\begin{aligned} & \underset{\{\mathbf{W}_i\}, \{\mathbf{b}_i\}}{\text{minimize}} && \mathcal{F}(\{\mathbf{W}_i\}_{i=1}^N, \{\mathbf{b}_i\}_{i=1}^N), \\ & \text{subject to} && \mathbf{W}_i \in \mathbf{S}_i, i = \{1, \dots, N\} \end{aligned} \quad (1)$$

where $\mathbf{S}_i = \{\text{remaining non-zeros weights in convolutional layer satisfy the requirement of structured pruning}\}$ is the pruning constraint set. We use the method proposed in ADMM-NN [15] to solve the problem.

B. Accelerator Enabled Embedded Software (ACE)

To accelerate the on-device DNN inference for the pruned DNN model, we propose vector accelerator-enabled embedded software (*ACE*), which deploys DNN algorithms efficiently and ensures full utilization of the vector operations provided by the accelerators, such as FFT, IFFT, MAC, etc. Instead of element-wise computation, *ACE* restructures the DNN computation by approaching bulk movement and vector operations in accelerators for both convolution and fully connected layer. Moreover, accelerator-based computation is susceptible to data overflow and executes only with fixed-point quantization. Therefore, *ACE* performs overflow-aware computation with 16 bit quantization. In addition, *ACE* proposes circular buffer convolution technique to reduce the memory footprint where the input/output memory is reused continuously.

Acceleration-aware dataflow: Figure 3 demonstrates the on-device data flow of LeNet-5 model as an example. Before executing each layer, the input and kernel are buffered in the SRAM allocated specifically for the accelerators. The buffered vectors are then sent to the accelerator for computation. Next, the outputs from the accelerators, buffered in SRAM, are stored in the FRAM. The subsequent Maxpooling layer and activation function (ReLU) is loaded directly into the CPU without SRAM buffering. *ACE* also selects the right kind of data movement method based on the energy and latency of moving the data. For example, large vector of data is moved with DMA while a single data is moved with CPU. During the inference, more time is spent on data movement, thus utilizing DMA with bulk data transfer achieves significant improvement over CPU-based data transfer.

Quantization: *ACE* provides software support for 16-bit fixed-point quantization to represent the floating-point data. $B \Rightarrow A * 2^{b-1}$ is the common quantization rule for embedded devices where "A" is a floating-point number, "B" is a quantized number, and "b" represents quantization bit.

Hardware Acceleration of FC layer: With the FFT/IFFT low energy vector operations, the DNN inference time and energy can be reduced significantly for BCM-based fully connected




Fig. 3: An example of on-device DNN data flow using LeNet-5.

Algorithm 1 On-device BCM implementation

```

1: Input:  $W, I$ 
2: Output:  $O$ 
3:  $I \leftarrow \text{SCALE-DOWN}(I, \text{len}(I))$ 
4:  $W \leftarrow \text{SCALE-DOWN}(W, \text{len}(W))$ 
5:  $cI \leftarrow \text{COMPLEX}(I)$ 
6:  $cW \leftarrow \text{COMPLEX}(W)$ 
7:  $cOut \leftarrow \text{IFFT}(\text{FFT}(cI) * \text{FFT}(cW))$ 
8:  $O \leftarrow \text{REAL}(cOut)$ 
9:  $O \leftarrow \text{SCALE-UP}(O, \text{len}(I), \text{len}(W))$ 
10: return  $O$ 
11: procedure SCALE-DOWN( $D, \text{length}$ )
12:   for each element  $\in \mathcal{D}$  do
13:     element  $\leftarrow$  element/length
14:   end for
15:   return  $\mathcal{D}$ 
16: end procedure
17: procedure SCALE-UP( $D, lI, lW$ )
18:   for each element  $\in \mathcal{D}$  do
19:     element  $\leftarrow$  element *  $lI * lW$ 
20:   end for
21:   return  $\mathcal{D}$ 
22: end procedure

```

(FC) layers. We redesign the computation of FC layer as shown in Algorithm 1, which describes the on-device implementation of BCM computation. In lines 3 and 4, input and weight data are scaled-down to overcome FFT operation overflow. Since data is not in complex form during the convolution, the real input needs to be converted into the complex number before performing FFT operation. ACE performs complex/real operations in lines 5-8. Line 7 performs the core function of the BCM based DNN. It performs FFT, complex number multiplication on both inputs, and then IFFT on the result. After that, real number is retrieved at line number 8 and stored to FRAM. In Line 9, output data is scaled up.

Hardware Acceleration of CONV layer: Instead of element-wise computation, ACE implements *bulk computation* by computing the whole kernel at a time. For example, as Figure 4, a typical convolution of a 3x3 kernel needs 18 operations (9 multiplication and 9 addition) with an input window. Instead, an accelerator can replace the whole with a single MAC [17] operation, resulting in drastic improvement in overall convolution performance.

Overflow-aware Computation: Fixed point calculation in resource constraint devices frequently suffers from data overflow errors. During FFT computation, there is a possible data overflow if the totaling of the input array exceeds the capacity of the quantized bit. For example, in an 8-bit data quantization, the FFT will produce wrong results if the addition of the input array elements exceeds $2^8 = 256$. To address the challenge, ACE




Fig. 4: Example of MAC Operation

performs data scaling based on the array size in Algorithm 1. Besides, there is also possible overflow during accelerator operation. The solution is provided using normalization as discussed in Section III-A.




Fig. 5: Circular buffer for input/output buffer reuse.

Circular Buffer Convolution: ACE optimizes the memory usage by reusing the input/output buffer after a layer-level computation takes place. In the existing DNN implementation, an inference with N layers requires N buffers, as shown in figure 5. Instead of allocating memory for individual layers, ACE requires only two buffers (input and output) at most. ACE implemented circular buffer convolution, restructured the DNN inference, which reuses the buffer by interchanging and overwriting the input and output pointer after finishing a layer-level computation. ACE can considerably reduce the memory footprint, irrespective of the layer size in a DNN inference. The size required for the buffer is $\max(L_i), 1 \leq i \leq N$.

C. Intermittent Inference with FLEX.

We designed FLEX, a software support that enables intermittent inference when there are frequent power failures. Previous work can successfully checkpoint the system state, which, however, does not work efficiently due to the large amount of intermediate results and computation state that takes much time and energy. A recent work [7] proposes TAILS that exploits the special structured and loop-heavy computations of DNN and enables their intermittent executions by continuously saving the loop control states. However, it will cause progress setbacks when FFT based computation is introduced.

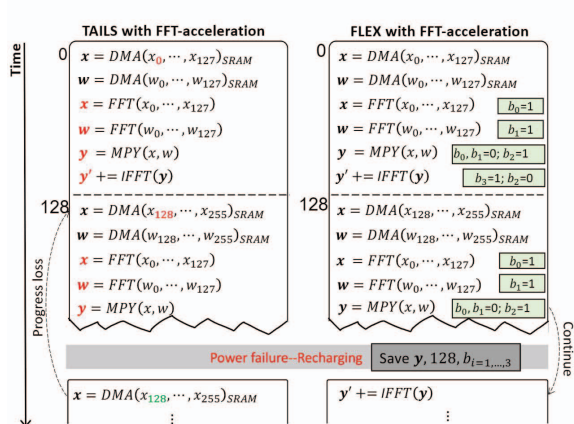


Fig. 6: Comparison between TAILS and FLEX.

BCM-based FC layer: The left part of Figure 6 demonstrates the drawback of *TAILS* when we try to employ FFT-based BCM computation under the intermittent power supply. An FFT-based BCM computation has to go through several steps including DMA, FFT/IFFT, and MPY as shown in algorithm 1. The arrays highlighted in red, x, w, y, y' are intermediate values. In loop-index based techniques, these are very much susceptible to data loss whenever a power failure happens. On power restore, the program has to roll back to the initial *DMA* operation because only the loop index is insufficient to retrieve the current state of the Algorithm. It causes wasted work and hurts the progress. As a solution, *FLEX* stores block index, intermediate result and state bit of the Algorithm. As shown in the right part of Figure 6, *FLEX* specifies the 4 additional bits b_0 - b_2 to indicate the current state, resulting in a successful continuation from the interrupted operation on power failure. As the control bits are small, it requires small memory footprint and read/write energy cost.

Other layer: In BCM-based FC layers, with the help of a voltage monitor system, *FLEX* predicts a power failure and checkpoint the latest intermediate result. Apart from that, for all other layers, *FLEX* follows loop index-based checkpointing mechanism that checkpoint and restores the system, based on loop indices.

D. Experimental Setup

Hardware Setup: The proposed DNN implementation framework is evaluated with TI's MSP430FR5994 [18] ultra-low-power evaluation board, consisting of a 16 MHz MCU, a 8KB volatile SRAM, a 256KB nonvolatile FRAM memory, and a low-energy Accelerator (LEA) that runs independently of CPU. The LEA accelerator supports FFT, IFFT, MAC, and MPY operation. The board is connected to the function generator SIGLENT SDG1032X [1] to simulate the energy harvesting scenario. Energy is buffered with a capacitor of 100 μ F. For the energy measurement, we used CCS energy trace technology [16].

TABLE II: Structure and Accuracy of DNN

Tasks	Layer	Original Size	Compress Method	Compression	Accuracy
MNIST	Conv	6 x 1 x 5 x 5	—	—	99%
	Conv	16 x 6 x 5 x 5	Structured Pruning	2x	
	FC	256 x 256	BCM	128x	
	FC	256 x 10	—	—	
HAR	Conv	32 x 1 x 1 x 12	—	—	89%
	FC	3520 x 128	BCM	128x	
	FC	128 x 64	BCM	64x	
OKG	Conv	6 x 1 x 5 x 5	—	—	82%
	FC	3456 x 512	BCM	256x	
	FC	512 x 256	BCM	128x	
	FC	256 x 128	BCM	64x	
	FC	128 x 12	—	—	

IV. EXPERIMENTS

DNN Models: This paper considers three DNN models for Image Classification (MNIST) [10], Human Activity Recognition (HAR) [2], and Google Keyword Recognition (OKG) [19] which represent image-based applications, wearable applications, and audio applications respectively as shown in Table II.

A. Experimental Results

We evaluated our framework by comparing with BASE, SONIC and TAILS [7], the state-of-the-art solutions. Here, BASE is the baseline implementation that does not consider intermittent operation.

1) *DNN Model Training and Pruning:* As shown in Table II, the structured pruning reduces the number of weights for the CONV layer by 2x and the BCM reduces the weights of the FC layer by 128x for the DNN model on MNIST task. For the HAR task, BCM reduces the number of weights by 128x for the first FC layer and the number of weights by 64x for the second FC layer. More FC layers of DNN models on the OKG dataset are compressed via BCM. The weights of the first three FC layers on this model are reduced for for 256x, 128x and 64x respectively.

2) *Inference Time under Continuous Power Supply:* We can observe from the figure 7(a) that *ACE* and *FLEX* run 3x, 4x, 3.3x faster inference than the Base, SONIC, and TAILS, respectively, on MNIST dataset. For HAR dataset, it is 5.4x, 5.7x, 2.6x faster. And lastly, it performs 1.7x, 3.3x, 2.1x faster on OKG dataset. This significant performance improvement is achieved by allowing the on-board acceleration engine to perform at the full extent with DMA-based data transfer. During the inference, we can see that most of the computation time is spent on convolutional layer while FC layer runs extremely fast.

3) *Inference Time under Intermittent Power Supply:* Under intermittent power supply, as shown in Figure 7(b), the Base model and *ACE* can never be completed because they do not tolerate intermittence. However, *ACE* with *FLEX* can successfully complete the inference because *FLEX* is developed to provide intermittent support. Besides, *FLEX* also reduced the wasted work and checkpointing overhead. Due to power failure, there is a negligible increase (1%-2%) in latency and energy consumption, achieving almost similar latency and energy as continuous power. Here, *ACE* + *FLEX* run 5.1x, 3.8x faster than the SONIC, and TAILS, respectively on MNIST dataset. For HAR dataset, it is 4.7x, 2.4x faster. And lastly, it performs 3.3x, 1.7x faster on OKG.




Fig. 7: Comparison of DNN Implementation with Different Framework




Fig. 8: Performance of First FC of MNIST.

4) *Energy Consumption and Performance*: In terms of energy consumption, our techniques performed ever better. We can observe from figure 7(c) that *ACE* and *FLEX* outperform *SONIC* and *TAILS* by achieving 6.1x and 4.31x energy-saving on *MNIST*. 10.9x and 5.26x energy-saving on *HAR*. And finally, 6.25x and 3.05x energy is saved on *OKG*. *LEA* and *DMA* run in ultra-low power mode and resulting in a significant amount of energy saving.

Specifically, Figure 8 demonstrates how *ACE* is successful in improving the overall performance of an FC layer. *ACE* is performed both independently and with different *BCM* block sizes (32, 64, 128) on the first FC layer of *MNIST* model. The employment of *BCM* based DNN inference proposed in this paper can significantly reduce the latency and energy needed. A larger block size of *BCM* reveals better performance and more compression. However, selecting a larger block size is limited by device support and accuracy degradation, which proves the efficiency of our resource-constrained framework.

5) *Evaluation of checkpointing overhead*: The proposed checkpointing mechanism considers special structure of DNN models and thus only saves latest intermediate result, necessary loop indices, and 4 control bits as shown in Figure 6. Every checkpoint/restore cost is at most 0.033mj, which is reached if power failure happens when computing the FFT-based *BCM* in the FC layer. Therefore, in our experiment, the total checkpoint/restore overhead is only 1%, 1.25%, and 0.8% for *MNIST*, *HAR*, and *OKG* datasets, respectively, which is negligible considering the largely overall saved energy.

V. CONCLUSION

This paper proposes an efficient framework for DNN implementation on energy harvesting devices which includes a resource-aware DNN training and pruning method considering hardware resource constraints, a DNN implementation method utilizing accelerators, and software support for intermittent computation that's friendly to long vector operations with accelerators. Specifically, the *BCM* compression is successfully

implemented with the FFT vector operation accelerators and the proposed quantization method which can improve the performance of FC layers by tens of times. The experimental results demonstrate significantly reduced DNN footprint, energy consumption, and improved performance.

REFERENCES

- [1] Sdg1032x. <https://siglentna.com/product/sdg1032x/>.
- [2] D. Anguita et al. A public domain dataset for human activity recognition using smartphones. In *Esann*, volume 3, page 3, 2013.
- [3] S. Boyd et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [4] S. Choi et al. An optimized design technique of low-bit neural network training for personalization on iot devices. In *DAC*, pages 1–6, 2019.
- [5] C. Ding et al. CirCNN: Accelerating and Compressing Deep Neural Networks using Block-circulant Weight Matrices. In *MICRO*, 2017.
- [6] Y. Geng et al. Tinyadc: Peripheral circuit-aware weight pruning framework for mixed-signal dnn accelerators. In *DATE 2021*. IEEE.
- [7] G. Gobieski et al. Intelligence beyond the edge: Inference on intermittent embedded systems. In *ASPLOS*, pages 199–213, 2019.
- [8] B. Jishnu et al. Memory-aware efficient deep learning mechanism for iot devices. In *ASAP*, 2021.
- [9] C. Kang et al. Everything leaves footprints: Hardware accelerated intermittent deep inference. In *ESWEEK-TCAD*, pages 3479–3491, 2020.
- [10] Y. LeCun and C. Cortes. *MNIST handwritten digit database*. 2010.
- [11] S. Lee and S. Nirjon. Neuro. zero: a zero-energy neural network accelerator for embedded sensing and inference systems. In *ENSS*, 2019.
- [12] C. Luo et al. Cosine normalization: Using cosine similarity instead of dot product in neural networks. In *ICANN*, pages 382–391, 2018.
- [13] H. Mendis et al. Intermittent-aware neural architecture search. In *ESWEEK-TECS*, pages 1–27, 2021.
- [14] A. Montanari et al. eperceptive—energy reactive embedded intelligence for batteryless sensors. In *SenSys*, 2020.
- [15] A. Ren et al. Admm-nn: An algorithm-hardware co-design framework of dnns using alternating direction methods of multipliers. In *ASPLOS 2019*, pages 925–938.
- [16] TI. Energy trace. <https://www.ti.com/tool/ENERGYTRACE>.
- [17] TI. Low energy accelerator. <http://www.ti.com/lit/an/slaa720/slaa720.pdf>.
- [18] TI. Msp430fr5994. <https://www.ti.com/product/MSP430FR5994>.
- [19] P. Warden. Speech commands: A dataset for limited-vocabulary speech recognition. 2018.
- [20] G. Yuan et al. A dnn compression framework for sot-mram-based processing-in-memory engine. In *SOCC 2020*, pages 37–42. IEEE.
- [21] G. Yuan et al. An ultra-efficient memristor-based dnn framework with structured weight pruning and quantization using admm. *ISLPED 2019*.
- [22] T. Zhang et al. A unified dnn weight pruning framework using reweighted optimization methods. In *DAC 2021*, pages 493–498.
- [23] S. Zhou et al. An end-to-end multi-task object detection using embedded gpu in autonomous driving. In *ISQED 2021*, pages 122–128.