FFT-Based Deep Learning Deployment in Embedded Systems

Sheng Lin[1*], Ning Liu[1*], Mahdi Nazemi[2], Hongjia Li[1], Caiwen Ding[1], Yanzhi Wang[1], and Massoud Pedram[2]

[1]Dept. of Electrical Engineering & Computer Science, Syracuse University, Syracuse, NY, USA
[2]Dept. of Electrical Engineering, University of Southern California, Los Angeles, CA, USA
[1]{shlin,nliu03,hli42,cading,ywang393}@syr.edu,[2]{mnazemi,pedram}@usc.edu

Abstract—Deep learning has delivered its powerfulness in many application domains, especially in image and speech recognition. As the backbone of deep learning, deep neural networks (DNNs) consist of multiple layers of various types with hundreds to thousands of neurons. Embedded platforms are now becoming essential for deep learning deployment due to their portability, versatility, and energy efficiency. The large model size of DNNs, while providing excellent accuracy, also burdens the embedded platforms with intensive computation and storage. Researchers have investigated on reducing DNN model size with negligible accuracy loss. This work proposes a Fast Fourier Transform (FFT)-based DNN training and inference model suitable for embedded platforms with reduced asymptotic complexity of both computation and storage, making our approach distinguished from existing approaches. We develop the training and inference algorithms based on FFT as the computing kernel and deploy the FFT-based inference model on embedded platforms achieving extraordinary processing speed.

I. INTRODUCTION

Recently deep learning has outstood from traditional machine learning techniques in many application areas, especially in image and speech recognition [1], [2]. The excellence of deep learning has also resulted in explorations of several emerging real-world applications, such as self-driving systems [3], automatic machine translations [4], drug discovery and toxicology [5]. The deep learning is based on the structure of deep neural networks (DNNs), which consist of multiple layers of various types and hundreds to thousands of neurons in each layer. Recent evidence has revealed that the network depth is of crucial importance to the success of deep learning, and many deep learning models for the challenging ImageNet dataset are sixteen to thirty layers deep [1]. Deep learning achieves significant improvement in overall accuracy by extracting complex and high-level features at the cost of considerable up-scaling in the model size.

In the big data era and driven by the development of semiconductor technology, embedded systems are now becoming an essential computing platform with ever-increasing functionalities. At the same time, researchers around the world from both academia and industry have devoted significant efforts and resources to investigate, improve, and promote the applications of deep learning in embedded systems [6]. Despite the advantages in DNN recognition accuracy, the deep layered structure and large model size of DNNs also increase computational complexity and memory requirement. Researchers are faced with the following challenges when deploying deep learning models on embedded systems: (i) Confined by the communication bandwidth of embedded systems, which are usually mobile terminals, it is still challenging to download large-size DNN models, even which can be offline-trained in data centers. (ii) The large model size of deep learning also imposes stringent requirements on the computing resources and memory size of embedded systems.

Motivated by these challenges, it is intuitive to implement a reduced-size deep learning model with negligible accuracy loss. In fact, the state-of-the-art DNNs are often over-parameterized, hence the removal of redundant parameters in the deep learning models, if performed properly, will produce similar overall accuracy as the original models [1]. Encouraged by this discovery, various deep learning model compression approaches have been investigated [6]–[10], including weight precision reduction, network pruning, weight matrix factorization, etc. In this work, we propose a Fast Fourier Transform (FFT)-based DNN training and inference model suitable for embedded systems due to reduced asymptotic complexity of both computation and storage. Our approach has obvious advantages over existing works on deep learning model compression e.g., [6], [8], [9] in that those approaches result in an irregular network architecture that increases training and inference computation time, while our approach facilitates computation. Please also note the our proposed framework is distinct from the prior work of using FFT for convolutional layer acceleration by LeCun et al. [11], because this prior work can only achieve convolutional layer acceleration instead of simultaneous compression. We develop the training and inference algorithms based on FFT as the computing kernel and deploy the FFT-based inference model on embedded platforms. Experimental test results demonstrate that our model provides the optimization in different languages and achieve a significant improvement.

II. RELATED WORK

Over the past decade, a substantial number of techniques and strategies have been proposed to compress neural network size. Weight pruning [6] is a well-known effective approach, in which many weights with values of 0 are pruned to achieve high compression ratio. Other techniques such as threshold setting [6], biased weight decay [9], etc., could be integrated to the weight pruning procedure. Another simple and popular approach to DNN model compression is the low-rank approximation of the weight matrix [12]. To overcome the potential high accuracy loss after low-rank approximation, [13] proposed to perform fine-tuning for the post-factorization of low-rank weight matrices to retain accuracy . Lowering the presentation precision of weights is also an straightforward technique to reduce both the model size and computation cost of DNNs. A fixed-point implementation was explored to replace the original floating-point models [14]. Furthermore, designs with ultra-low precision weights, such as binary (-1 / +1) or ternary (-1 / 0 / +1) representation were proposed [15], [16]. By exploring the local and global characteristics of the weight matrix, weight clustering was proposed to reduce the number of weights linearly [17]. In addition, with the aid of gradients clustering in the training phase, the accuracy loss incurred by the weight clustering can be negligible [6].

Some recent works adopted structured weight matrices in

---

order to reduce the model size. In [18], weight matrices of fully-connected (FC) layers were constructed in the Toeplitz-like format to remove the redundancy of the DNN model. In [19], the circulant matrix was introduced to enable further reduction in model size. An n-by-n circulant matrix has a smaller number of parameters i.e., n than that of a same-size Toeplitz matrix i.e., $2n$. In this work, we generalize the structured weight matrix method in that (1) we utilize block-circulant matrices for weight matrix representation, which achieves a trade-off between compression ratio and accuracy loss; (2) we extend the structured matrix method to convolutional (CONV) layers besides the FC layers; (3) we propose FFT-based DNN training and inference model and algorithm, which is highly suitable for deployment in embedded systems; and (4) we implement and test the FFT-based DNN inference in various embedded platforms.

III. BACKGROUND

In this section, we introduce basic concepts of deep neural networks (DNNs), Fast Fourier Transform (FFT), and structured matrices, as the background of our proposed FFT-based training and inference algorithms. Specifically, we explain the various DNN layer types, the Cooley-Tukey algorithm for FFT, and the block-circulant matrices as the adopted structured matrices.

A. Deep Neural Networks

Deep neural networks (DNNs) are distinguished from other types of neural networks by their depth and have dramatically improved the state-of-the-art in speech recognition, object detection, etc. Some commonly adopted DNN models include deep convolutional neural networks, deep belief networks, and recurrent neural networks. Despite the various network topologies targeting for different applications, these DNN models comprise of multiple functional layers with some commonly used structures. Following are the most commonly used layer structures in the state-of-the-art DNN models:

The fully-connected (FC) layer is the most storage-intensive layer in DNN models [20] since each of its neurons is fully connected with all the neurons in the previous layer. The computation procedure of a FC layer consists of matrix-vector arithmetics (multiplication and addition) and transformation by the activation function, described as follows:

$$\mathbf{y} = \psi(\mathbf{W^T x} + \theta) \tag{1}$$

where $\mathbf{y}$ and $\mathbf{x}$ are outputs of this layer and the previous layer, respectively; $\mathbf{W} \in \mathbb{R}^{m \times n}$ is the weight matrix of the synapses between this FC layer (with n neurons) and its previous layer (with m neurons); $\theta \in \mathbb{R}^n$ is the bias vector; and $\psi(\cdot)$ is the activation function. The Rectified Linear Unit (ReLU) $\psi(x) = \max(0, x)$ is the most widely utilized activation function in DNNs.

The convolutional (CONV) layer, as the name implies, performs two-dimensional convolution of its input to extract features that will be fed into subsequent layers for higher-level feature extracting. A CONV layer is associated with a set of learnable filters [21], which are activated when specific types of features are found at some spatial positions from the inputs. Filter-sized moving windows are applied to the inputs to obtain a set of feature maps, by calculating the convolution of the filter and inputs in the moving window. Each *convolutional neuron*, representing one pixel in a feature map, takes a set of inputs and the corresponding filter weights to calculate the inner-product. Given input feature map $\mathbf{X}$ and the $r \times r$-sized
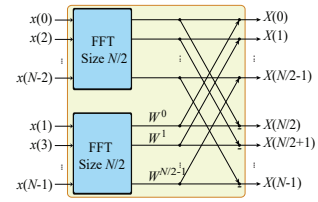

Fig. 1. Illustration of Cooley-Tukey algorithm of FFT.

filter (i.e., the *convolutional kernel*) $\mathbf{F}$, the output feature map $\mathbf{Y}$ is calculated as

$$y_{a,b} = \sum_{i=1}^{r} \sum_{j=1}^{r} x_{a+i-1, b+j-1} \times f_{i,j}, \tag{2}$$

where $y_{a,b}$, $x_{a+i-1, b+j-1}$, and $f_{i,j}$ are elements in $\mathbf{Y}$, $\mathbf{X}$, and $\mathbf{F}$, respectively. Multiple convolutional kernels can be adopted to extract different features in the same input feature map. Multiple input feature maps can be convolved with the same filter and results are summed up to derive a single feature map.

B. Fast Fourier Transforms

The Fast Fourier Transform (FFT) is an efficient procedure for computing the discrete Fourier transform (DFT) of time series. It takes advantage of the fact that the calculation of the coefficients of the DFT can be carried out iteratively, which results in a considerable savings of computation time. The FFT not only reduces the computational complexity, but also substantially reduces round-off errors associated with these computations. In fact, both the computation time and round-off error are essentially reduced by a factor of $n/(log_2 n)$ where n is the number of data samples in the time series [22]. Fig. 1 shows the simplest and most common form of FFT, which is based on the Cooley-Tukey algorithm [23]. It uses a divide and conquer approach to recursively break down the DFT of an arbitrary composite size $N = N_1 \cdot N_2$ into many smaller DFTs of sizes N_1 and N_2, in order to reduce the computation time to $\mathrm{O}(n \log n)$ for highly composite N [23].

C. Structured Matrices

An n-by-m matrix $\mathbf{A}$ is called a structured matrix when it has a low displacement rank υ [18]. One of the most important characteristics of structured matrices is their low number of independent variables. The number of independent parameters is $O(n)$ for an n-by-n structured matrix instead of $O(n^2)$, which indicates that the storage complexity can be potentially reduced to $O(n)$. As a representative example, a circulant matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ is defined by a vector $\mathbf{w} = (w_1, w_2, ..., w_n)$ as follows:

$$\begin{bmatrix} w_1 & w_n & \dots & w_3 & w_2 \\ w_2 & w_1 & w_n & & w_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ w_{n-1} & & \ddots & \ddots & w_n \\ w_n & w_{n-1} & \dots & w_2 & w_1 \end{bmatrix}$$

The definition and analysis of structured matrices have been generalized to the case of m-by-n matrices where $m \neq n$, e.g., the block-circulant matrices. Besides, the computational complexity for many matrix operations, such as matrix-vector multiplication, matrix inversion, etc., can be significantly reduced when operating on structured matrices.

IV. FAST FOURIER TRANSFORM-BASED DNN MODEL

In this section, we propose an efficient inference algorithm and explain the training algorithm in deep neural networks by using block-circulant matrices. We achieve a simultaneous and significant reduction in computational complexity of inference and training processes, and also weight storage. Besides, we have performed theoretical analysis to prove the effectiveness of substituting matrix multiplication with the Fast Fourier Transform method and utilizing block-circulant matrices, thereby guaranteeing applicability of the proposed framework on a wide variety of applications and emerging deep learning models.

A. Block-Circulant Matrix-Based Inference and Training Algorithms for FC Layers

Cheng et al. proposed circulant matrix-based DNN training and inference algorithms for FC layers [19]. However, in many practical applications such schemes cannot be directly used because: (1) It is very common that the weight matrices of DNNs are non-square matrices due to the specific need of different applications; and (2) Even if the weight matrices are square, in many cases the compression is too aggressive and hence causes non-negligible accuracy loss. To address the above challenges, we present the block-circulant matrix-based inference and training algorithms.

Recall that the forward propagation during the inference phase of a FC layer is performed as $\mathbf{y} = \psi(\mathbf{W}^{\mathbf{T}}\mathbf{x} + \theta)$, where ψ is the activation function, $\mathbf{W}$ is the weight matrix, $\mathbf{x}$ is the input vector, and θ is the biases. The computation bottleneck is the calculation of $\mathbf{W}^{\mathbf{T}}\mathbf{x}$. When using a block-circulant matrix for representing $\mathbf{W}$, a fast multiplication algorithm for $\mathbf{W}^{\mathbf{T}}\mathbf{x}$ exists, which will result in a significant reduction in computational complexity. Assume that the weight matrix is an m-by-n block-circulant matrix $\mathbf{W} = [\mathbf{C}_1|\mathbf{C}_2|...|\mathbf{C}_k]^{\mathbf{T}}$; the input vector is $\mathbf{x} = (\mathbf{x}_1|\mathbf{x}_2|...|\mathbf{x}_k)$; and the bias vector is $\theta = (\theta_1|\theta_2|...|\theta_k)$. Each circulant matrix $\mathbf{C}_i \in \mathbb{R}^{n \times n}$ is defined by a length-n vector $\mathbf{w}_i = (w_{i,1}, w_{i,2}, ..., w_{i,n})^{\mathbf{T}}$, $i \in \{1,...,k\}$, $m = kn$, and $\mathbf{x}_i = (x_{i,1}, x_{i,2}, ..., x_{i,n})^{\mathbf{T}}$. Hence, $\mathbf{W}^{\mathbf{T}}\mathbf{x}$, as the key computation bottleneck in the inference phase, can be simplified as below:

$$\mathbf{W}^{\mathbf{T}}\mathbf{x} = \sum_{i=1}^{k} \mathbf{C}_i \mathbf{x}_i = \sum_{i=1}^{k} \text{IFFT}\big(\text{FFT}(\mathbf{w}_i) \circ \text{FFT}(\mathbf{x}_i)\big) \quad (3)$$

where FFT, IFFT, and $\circ$ represent a Fast Fourier transform (FFT), an inverse FFT, and an element wise multiplication, respectively. This "FFT $\rightarrow$ component-wise multiplication $\rightarrow$ IFFT" procedure to implement $\mathbf{W}^{\mathbf{T}}\mathbf{x}$ shown in Fig. 2 is derived from the circular convolution theorem [24], [25]. The overall computational complexity in this FC layer will be O($n \log n$), achieving a significant reduction compared to O(n^2) when calculating $\mathbf{W}^{\mathbf{T}}\mathbf{x}$ directly. In order to store the weights for the inference phase, we can simply keep the FFT result $\text{FFT}(\mathbf{w}_i)$ (which is a vector) instead of the whole matrix $\mathbf{W}$, thereby reducing the storage complexity to O(n) for an FC layer. Algorithm 1 summarizes the FFT-based inference algorithm.

Besides the inference procedure, the reformulated training (weight updating) algorithm in the scenario of using block-circulant matrices will also result in significant accelerations. We denote $\mathbf{a} = \mathbf{W}^{\mathbf{T}}\mathbf{x} + \theta = (\mathbf{a}_1|\mathbf{a}_2|...|\mathbf{a}_k)^{\mathbf{T}}$ and $\mathbf{a}_i =$

Algorithm 1: Block-circulant Matrix-based Inference

Input: w, x, m, n
Output: a
$sa \leftarrow max(m, n)$;
$si \leftarrow min(m, n)$;
$k \leftarrow \lceil sa/si \rceil$;
partition w into k vectors, $w_1, \ldots, w_k$;
if $m > n$ **then**
    for $i \leftarrow 0$ *until* k **do**
        $a \leftarrow a + \text{ifft}(\text{fft}(\mathbf{w_i}) \circ \text{fft}(x))$;
    end
else
    partition x into k vectors, $x_1, \ldots, x_k$;
    for $i \leftarrow 0$ *until* k **do**
        $a \leftarrow a + \text{ifft}(\text{fft}(\mathbf{w_i}) \circ \text{fft}(x_i))$;
    end
end
return a;

$(a_{i,1}, a_{i,2}, ..., a_{i,n})^{\mathbf{T}}$, then the weight updating rule for the block-circulant FC layer is given by:

$$\mathbf{w}_i \leftarrow \mathbf{w}_i - \epsilon \cdot \text{IFFT}\Big(\text{FFT}\big(\frac{\partial L}{\partial \mathbf{a}_i}\big) \circ \text{FFT}(\mathbf{x}'_i)\Big) \cdot \mathbf{I} \quad (4)$$

where J, $\mathbf{I}$, ϵ, and $\mathbf{x}'_i$ represent the loss function, an all-one column vector, the learning rate, and the base vector that defines the circulant matrix $\frac{\partial \mathbf{a}_i}{\partial \mathbf{w}_i}$ (which is formally derived), respectively. Notice that since $\frac{\partial \mathbf{a}_i}{\partial \mathbf{w}_i}$ is a circulant matrix, similar to inference, we can utilize the "FFT→component-wise multiplication→IFFT" procedure to accelerate the matrix-vector multiplication. The computational complexity will be O($n \log n$) in each updating step in this layer, which is a significant reduction from O(n^2) in traditional backpropagation procedure. Algorithm 2 summarizes the FFT-based training algorithm.

B. Block-Circulant Matrix-Based Inference and Training Algorithms for CONV Layer

The use of block-circulant matrices can also enable significant reduction in computational and storage complexities of the Convolutional layer. The Convolutional layers are often associated with multiple input and output feature maps in DNNs. Therefore, the computation of the Convolutional layer is described using tensor format as follows:

$$\mathcal{Y}(x, y, p) = \sum_{i=1}^{r} \sum_{j=1}^{r} \sum_{c=1}^{C} \mathcal{F}(i, j, c, p) \mathcal{X}(x+i-1, y+j-1, c),$$

$$(5)$$

where $\mathcal{X} \in \mathbb{R}^{W \times H \times C}$, $\mathcal{Y} \in \mathbb{R}^{(W-r+1) \times (H-r+1) \times P}$, $\mathcal{F} \in \mathbb{R}^{r \times r \times C \times P}$ denote the input, output, and weight "tensors" of
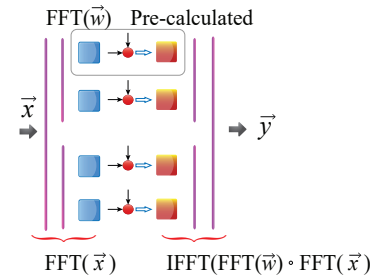

Fig. 2. The "FFT $\rightarrow$ component-wise multiplication $\rightarrow$ IFFT" procedure.

For general values of m and n, we can apply zero padding such that the definition of block-circulant matrices can be applied.

Algorithm 2: Block-circulant Matrix-based Training

Input: $\frac{\partial L}{\partial a}$, w, x, m, n
Output: $\frac{\partial L}{\partial w}$, $\frac{\partial L}{\partial x}$
$sa \leftarrow max(m, n)$;
$si \leftarrow min(m, n)$;
$k \leftarrow \lceil sa/si \rceil$;
partition w into k vectors, $w_1, \ldots, w_k$;
partition $\frac{\partial L}{\partial w}$ into k vectors, $\frac{\partial L}{\partial w_1}, \ldots, \frac{\partial L}{\partial w_k}$;
if $m > n$ **then**
    partition $\frac{\partial L}{\partial a}$ into k vectors, $\frac{\partial L}{\partial a_1}, \ldots, \frac{\partial L}{\partial a_k}$;
    for $i \leftarrow 0$ *until* k **do**
        $\frac{\partial L}{\partial w_i} \leftarrow \text{ifft}(\text{fft}(\frac{\partial L}{\partial a}) \circ \text{fft}(\mathbf{x'})) \cdot \mathbf{1}$;
        $\frac{\partial L}{\partial x} \leftarrow \frac{\partial L}{\partial x} + \text{ifft}(\text{fft}(\frac{\partial L}{\partial a}) \circ \text{fft}(\mathbf{w_i'}))$;
    end
else
    partition x into k vectors, $x_1, \ldots, x_k$;
    partition $\frac{\partial L}{\partial x}$ into k vectors, $\frac{\partial L}{\partial x_1}, \ldots, \frac{\partial L}{\partial x_k}$;
    for $i \leftarrow 0$ *until* k **do**
        $\frac{\partial L}{\partial w_i} \leftarrow \text{ifft}(\text{fft}(\frac{\partial L}{\partial a}) \circ \text{fft}(\mathbf{x_i'})) \cdot \mathbf{1}$;
        $\frac{\partial L}{\partial x_i} \leftarrow \text{ifft}(\text{fft}(\frac{\partial L}{\partial a}) \circ \text{fft}(\mathbf{w_i'}))$;
    end
end
return $\frac{\partial L}{\partial w}$, $\frac{\partial L}{\partial x}$;

the Convolutional layer, correspondingly. C is the number of input maps. W and H are the spatial dimensions of the input maps. P is the total number of output maps, and r is the size of the convolutional kernel.

We generalize the "block-circulant structure" as rank-4 tensor ($\mathcal{F}$) in the Convolutional layer, i.e., *each slice $\mathcal{F}(\cdot, \cdot, i, j)$ is a circulant matrix*. Then, we reformulate the inference and training algorithms of the Convolutional layer to matrix-based operations.

In the Convolutional layer, to enhance the implementation efficiency, software tools provide an efficient approach of changing tensor-based operations to matrix-based operations equivalently [26], [27]. Fig. 3 demonstrates the application of the method to reformulate Eqn. (3) to the matrix multiplication $\mathbf{Y} = \mathbf{XF}$, where $\mathbf{X} \in \mathbb{R}^{(W-r+1)(H-r+1) \times Cr^2}$, $\mathbf{Y} \in \mathbb{R}^{(W-r+1)(H-r+1) \times P}$, and $\mathbf{F} \in \mathbb{R}^{Cr^2 \times P}$.
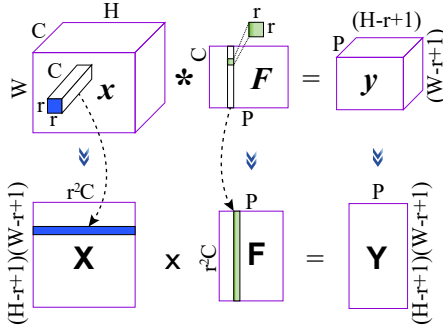

Fig. 3. Reformulation of Eqn. (3) to matrix multiplication.

Based on the reshaping principle between $\mathcal{F}$ and $\mathbf{F}$, we have:

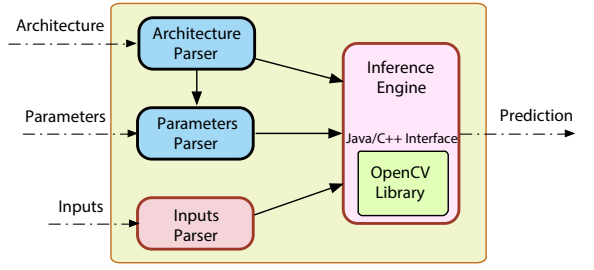$$f_{a+C(i-1)+Cr(j-1),b} = f_{C(i-1)+Cr(j-1),b-a}, \forall a, b \quad (6)$$

Fig. 4. Building blocks of software implementation.

where $\mathbf{F}$ is a block-circulant matrix. Therefore, the "FFT→component-wise multiplication →IFFT" procedure can be applied to accelerate $\mathbf{Y} = \mathbf{XF}$, leading to the acceleration of (3). With the assist of the proposed approach, the computational complexity for (3) is reduced from $O(WHr^2CP)$ to $O(WHQ \log Q)$, where $Q = \max(r^2C, P)$.

V. SOFTWARE IMPLEMENTATION

In this section, we provide detailed explanation of our software implementation, experimental setup, and evaluation of the proposed inference framework on various Android-based platforms with embedded processors and various datasets. The purpose of this software implementation is to reveal the potential of embedded systems in running real time applications that involve deep neural networks.

The software implementation of proposed inference framework for Android-based platforms is comprised of four high-level modules. The first module is responsible for constructing the network architecture. The second module reads a file that contains trained weights and biases. The third module loads test data that consists of input features and predefined classification labels, and finally, the fourth module performs inference for predicting labels. Fig. 4 depicts these high-level building blocks of the software implementation, along with their interactions. It should be noted that the test data may be loaded from a file, camera, etc.

We utilize the OpenCV[28] as core computing library in our project. OpenCV is an open-source cross-platform library of programming functions that is mainly targeted for computer vision applications and includes efficient implementation of aforementioned operations. OpenCV is written in C++, and it provides the API (Application Program Interface) for both C++ and Java. We implement two versions of software for inference: one that uses OpenCV's Java API, which is more convenient for Android development, and another one that is developed in C++ using Android NDK (Native Development Kit), uses OpenCV's C++ API, and is expected to have a better performance.

A. Experimental Setup

We run the inference application on various platforms of different generations in order to evaluate the applicability of the inference on embedded systems. Table I summarizes the specifications of test platforms.

The OpenCV Manager is installed on all target platforms in order to link OpenCV libraries dynamically and reduce memory usage. Additionally, hardware specific optimizations are applied by OpenCV Manager for an application's supported platforms.

TABLE I.   PLATFORMS UNDER TEST AND THEIR SPECIFICATIONS.

Platform	Android	Primary CPU	Companion CPU	CPU Architecture	GPU	RAM (GB)
LG Nexus 5	6 (Marshmallow)	4 × 2.3GHz Krait 400	-	ARMv7-A	Adreno 330	2
Odroid XU3	7 (Nougat)	4 × 2.1GHz Cortex-A15	4 × 1.5GHz Cortex-A7	ARMv7-A	Mali T628	2
Huawei Honor 6X	7 (Nougat)	4 × 2.1GHz Cortex-A53	4 × 1.7GHz Cortex-A53	ARMv8-A	Mali T830	3

In order to standardize the evaluation process on all platforms, the airplane mode is switched on to eliminate telecommunication overhead; all other running applications are closed to ensure they do not affect runtime; and the device is plugged in to avoid performance throttling applied by a platform's governor. Though this is the standard setup, we will study the performance of inference process in situations where the device is running on its battery.

B. MNIST

MNIST dataset [29] is a handwritten digits dataset which includes 28×28 greyscale images with 60,000 images for training and 10,000 images for testing. The original images in the MNIST dataset are resized using a bilinear transformation, and such transformation is used for both training and testing. Various neural network architectures are explored for each dataset and a few of them are presented in this paper.

For the MNIST dataset, two different neural network architectures are evaluated. In the first architecture (Arch. 1), the input layer consists of 256 neurons that represent the resized MNIST images. The next two layers comprise of 128 neurons each and are based on block-circulant matrix based FC layers. Finally, the last layer is a softmax layer that consists of 10 neurons representing the ten possible predictions for the digits. The second architecture (Arch. 2) has 121 neurons in the input layer, 64 neurons in the two hidden layers, and similar to Arch. 1, a softmax layer as the output layer. Table II summarizes the runtime of each round of inference process using these architectures and on various mobile platforms.

TABLE II.   CORE RUNTIME OF EACH ROUND OF INFERENCE FOR RESIZED MNIST IMAGES.

Architecture	Implementation	Accuracy (%)	Runtime (μs per image)		
			Nexus 5	XU3	Honor 6X
Arch. 1	Java	95.47	359.6	294.1	256.7
	C++	95.47	140.0	122.0	101.0
Arch. 2	Java	93.59	350.9	278.2	221.7
	C++	93.59	128.5	119.1	98.5

Based on the results summarized in Table II, the C++ implementation is about 60-65% faster than the Java implementation. One of the reasons for this superior performance is related to memory limitations and management policy in Android. While applications written in C++ have an unlimited heap size, Java applications are restricted to platform-specific heap sizes. As a result, a constraint is imposed on the amount of data that an application written in Java can deal with at each instance of time.

Another potential reason that may explain the considerable performance difference between the two implementations is the overhead due to switching from Java to C++ and vice versa. Because the OpenCV library is written in C++, it needs to covert data from C++ data types to Java data types whenever the Java API is used. We believe that these conversions do not affect the runtime significantly, but can cause certain difference in performance across the two implementations.

Considering different architectures mentioned in Table II, one can observe that going from the smaller network to a bigger network increases the accuracy by about 2% while it increases the memory required for storing parameters by a factor of about two and increases the runtime of Java and C++ implementations by about 2% and 9%, respectively. It should be noted that when the device is running on its battery, the runtime will increase by about 14% in the Java implementation, but remains unchanged in the C++ implementation.

C. CIFAR-10

The CIFAR-10 [30] dataset contains 32×32 color images from 10 classes, where there are 50,000 training images and 10,000 testing images. The structure of deep neural network can be denoted as 128x3x32x32-64Conv3-64Conv3-128Conv3-128Conv3-512F-1024F-1024F-10F (Arch. 3). Here 128x3x32x32 represents that (i) the batch size is 128; (ii) the number of input channel is 3, (iii) and the feature size of input data is 32x32. In addition, 128Conv3 indicates that 128 3x3 convolutional filters are used in the convolutional layer. In addition, 512F or 10F means that the number of neurons in the FC layer is 512 or 10, respectively. In addition, both the original and compressed models are trained with learning rate 0.001 and momentum 0.9. In this network architecture, the first two convolutional layers are traditional convolutional layers (no block circulant, which is treated as preprocessing similar to the IBM TrueNorth paper [31]). Based on the results summarized in Table III, the C++ implementation is about 130% faster than the Java implementation.

TABLE III.   CORE RUNTIME OF EACH ROUND OF INFERENCE PROCESS FOR CIFAR-10 IMAGES.

Architecture	Implementation	Accuracy (%)	Runtime (μs per image)	
			XU3	Honor 6X
Arch. 3	Java	80.2	21032	19785
	C++	80.2	8912	8244

D. Comparison Results on Performance and Accuracy

In this section, we provide comprehensive comparison results on MNIST, CIFAR-10, and IBM TrueNorth [31], [32]. Our test platform consists of one or two qual-core ARM, while the IBM TrueNorth includes 4,096 ASIC cores, which is around 500-1000 times more than our testing platform. In Fig. 5, compared with IBM TrueNorth results on MNIST [32], our model performs 10× faster than IBM TrueNorth with a little accuracy reduction on the best device result. The accuracy for IBM TrueNorth is 95% and the runtime is 1000μs per image on MNIST. Compared with IBM TrueNorth results on CIFAR-10 [31], with 500-1000 times less cores, our model performs 10× slower than IMB TrueNorth. The accuracy for IBM TrueNorth is 83.41% and the runtime is 800μs per image. We can see that the later work [31] in 2016 on CIFAR-10 is optimized more efficiently compared with the former work [32] in 2015. Although our mobile phone based framework achieves lower performance compared with IBM TrueNorth on CIFAR-10, it is still reasonably good result considering the
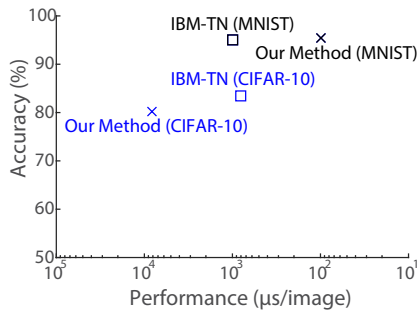
Fig. 5. Performance vs. accuracy results comparison on the MNIST and CIFAR-10 benchmarks.

dramatic difference in computational resources. These results have demonstrated the effectiveness of the proposed framework.

VI. CONCLUSIONS

This paper presented a design optimization framework for Fast Fourier Transform-based deep neural network inference on embedded system. The proposed approach results in significant reduction in storage requirement for model parameters and improves runtime without affecting accuracy significantly. Our implementation on ARM-based embedded systems achieves runtime improvement on image classification tasks compared to IBM TrueNorth.

VII. ACKNOWLEDGEMENT

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[2] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Acoustics, speech and signal processing (icassp), 2013 ieee international conference on*. IEEE, 2013, pp. 6645–6649.

[3] B. Huval, T. Wang, S. Tandon, J. Kiske, W. Song, J. Pazhayampallil, M. Andriluka, P. Rajpurkar, T. Migimatsu, R. Cheng-Yue *et al.*, "An empirical evaluation of deep learning on highway driving," *arXiv preprint arXiv:1504.01716*, 2015.

[4] R. Collobert and J. Weston, "A unified architecture for natural language processing: Deep neural networks with multitask learning," in *Proceedings of the 25th international conference on Machine learning*. ACM, 2008, pp. 160–167.

[5] R. Burbidge, M. Trotter, B. Buxton, and S. Holden, "Drug design by machine learning: support vector machines for pharmaceutical data analysis," *Computers & chemistry*, vol. 26, no. 1, pp. 5–14, 2001.

[6] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," *arXiv preprint arXiv:1510.00149*, 2015.

[7] A. Ren, Z. Li, C. Ding, Q. Qiu, Y. Wang, J. Li, X. Qian, and B. Yuan, "Sc-dcnn: Highly-scalable deep convolutional neural network using stochastic computing," in *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2017, pp. 405–418.

[8] Y. LeCun, J. S. Denker, S. A. Solla, R. E. Howard, and L. D. Jackel, "Optimal brain damage." in *NIPs*, vol. 2, 1989, pp. 598–605.

[9] L. Y. Pratt, *Comparing biases for minimal network construction with back-propagation*. Morgan Kaufmann Pub, 1989, vol. 1.

[10] M. Nazemi, S. Nazarian, and M. Pedram, "High-performance FPGA implementation of equivariant adaptive separation via independence algorithm for Independent Component Analysis," in *Application-specific Systems, Architectures and Processors (ASAP), 2017 IEEE 28th International Conference on*. IEEE, 2017, pp. 25–28.

[11] M. Mathieu, M. Henaff, and Y. LeCun, "Fast training of convolutional networks through ffts," *arXiv preprint arXiv:1312.5851*, 2013.

[12] M. Denil, B. Shakibi, L. Dinh, N. de Freitas *et al.*, "Predicting parameters in deep learning," in *Advances in Neural Information Processing Systems*, 2013, pp. 2148–2156.

[13] J. Chung and T. Shin, "Simplifying deep neural networks for neuromorphic architectures," in *Design Automation Conference (DAC), 2016 53nd ACM/EDAC/IEEE*. IEEE, 2016, pp. 1–6.

[14] S. Anwar, K. Hwang, and W. Sung, "Fixed point optimization of deep convolutional neural networks for object recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1131–1135.

[15] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1," *arXiv preprint arXiv:1602.02830*, 2016.

[16] K. Hwang and W. Sung, "Fixed-point feedforward deep neural network design using weights+ 1, 0, and- 1," in *Signal Processing Systems (SiPS), 2014 IEEE Workshop on*. IEEE, 2014, pp. 1–6.

[17] Y. Gong, L. Liu, M. Yang, and L. Bourdev, "Compressing deep convolutional networks using vector quantization," *arXiv preprint arXiv:1412.6115*, 2014.

[18] V. Sindhwani, T. Sainath, and S. Kumar, "Structured transforms for small-footprint deep learning," in *Advances in Neural Information Processing Systems*, 2015, pp. 3088–3096.

[19] Y. Cheng, F. X. Yu, R. S. Feris, S. Kumar, A. Choudhary, and S.-F. Chang, "An exploration of parameter redundancy in deep networks with circulant projections," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2857–2865.

[20] J. Qiu, J. Wang, S. Yao, K. Guo, B. Li, E. Zhou, J. Yu, T. Tang, N. Xu, S. Song *et al.*, "Going deeper with embedded fpga platform for convolutional neural network," in *Proceedings of the 2016 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 2016, pp. 26–35.

[21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

[22] W. T. Cochran, J. W. Cooley, D. L. Favin, H. D. Helms, R. A. Kaenel, W. W. Lang, G. Maling, D. E. Nelson, C. M. Rader, and P. D. Welch, "What is the fast fourier transform?" *Proceedings of the IEEE*, vol. 55, no. 10, pp. 1664–1674, 1967.

[23] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex fourier series," *Mathematics of computation*, vol. 19, no. 90, pp. 297–301, 1965.

[24] V. Pan, *Structured matrices and polynomials: unified superfast algorithms*. Springer Science & Business Media, 2012.

[25] L. Zhao, S. Liao, Y. Wang, Z. Li, J. Tang, and B. Yuan, "Theoretical properties for neural networks with weight matrices of low displacement rank," *international conference on machine learning*, pp. 4082–4090, 2017.

[26] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 2014, pp. 675–678.

[27] A. Vedaldi and K. Lenc, "Matconvnet: Convolutional neural networks for matlab," in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 689–692.

[28] Itseez, "Open source computer vision library," https://github.com/itseez/opencv, 2015.

[29] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010. [Online]. Available: http://yann.lecun.com/exdb/mnist/

[30] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009.

[31] S. K. Esser, P. A. Merolla, J. V. Arthur, A. S. Cassidy, R. Appuswamy, A. Andreopoulos, D. J. Berg, J. L. McKinstry, T. Melano, D. R. Barch *et al.*, "Convolutional networks for fast, energy-efficient neuromorphic computing," *Proceedings of the National Academy of Sciences*, p. 201604850, 2016.

[32] S. K. Esser, R. Appuswamy, P. Merolla, J. V. Arthur, and D. S. Modha, "Backpropagation for energy-efficient neuromorphic computing," in *Advances in Neural Information Processing Systems*, 2015, pp. 1117–1125.