

Elastic Significant Bit Quantization and Acceleration for Deep Neural Networks

Cheng Gong¹, Ye Lu¹, Kunpeng Xie¹, Zongming Jin, Tao Li¹, and Yanzhi Wang², *Member, IEEE*

Abstract—Quantization has been proven to be a vital method for improving the inference efficiency of deep neural networks (DNNs). However, it is still challenging to strike a good balance between accuracy and efficiency while quantizing DNN weights or activation values from high-precision formats to their quantized counterparts. We propose a new method called elastic significant bit quantization (ESB) that controls the number of significant bits of quantized values to obtain better inference accuracy with fewer resources. We design a unified mathematical formula to constrain the quantized values of the ESB with a flexible number of significant bits. We also introduce a distribution difference aligner (DDA) to quantitatively align the distributions between the full-precision weight or activation values and quantized values. Consequently, ESB is suitable for various bell-shaped distributions of weights and activation of DNNs, thus maintaining a high inference accuracy. Benefitting from fewer significant bits of quantized values, ESB can reduce the multiplication complexity. We implement ESB as an accelerator and quantitatively evaluate its efficiency on FPGAs. Extensive experimental results illustrate that ESB quantization consistently outperforms state-of-the-art methods and achieves average accuracy improvements of 4.78%, 1.92%, and 3.56% over AlexNet, ResNet18, and MobileNetV2, respectively. Furthermore, ESB as an accelerator can achieve 10.95 GOPS peak performance of 1k LUTs without DSPs on the Xilinx ZCU102 FPGA platform. Compared with CPU, GPU, and state-of-the-art accelerators on FPGAs, the ESB accelerator can improve the energy efficiency by up to 65 \times , 11 \times , and 26 \times , respectively.

Index Terms—DNN quantization, significant bits, fitting distribution, cheap projection, distribution aligner, FPGA accelerator

1 INTRODUCTION

DEEP neural networks (DNNs) usually involve hundreds of millions of trained parameters and floating-point operations [1], [2]. Thus, it is challenging for devices with limited hardware resources and constrained power budgets to use DNNs [3]. Quantization has been proven to be an effective method for improving the computing efficiency of DNNs [4], [5]. The method maps the full-precision floating-point weights or activation into low-precision representations

to reduce the memory footprint of model and resource usage of multiply-accumulate unit [6], [7], [8]. For example, the size of a 32-bit floating-point DNN model can be reduced by 32 \times by quantizing its weights into binaries.

However, low-bitwidth quantization can lead to large error, thus resulting in significant accuracy degradation [9], [10], [11]. The error is the difference of weights/activation before and after quantization which can be quantitatively measured by euclidean distance [10]. Although 8-bit quantization, such as QAT [12] and μ L2Q [8], can mitigate the accuracy degradation, use of a high bitwidth incurs comparatively large overheads. It is still challenging to strive for high accuracy when using a small number of quantizing bits.

Preliminary studies [13], [14], [15] imply that fitting weight distribution of DNNs can improve accuracy. Fitting distribution refers to adopting a high resolution, i.e., more quantized values, for the densely distributed data range, such as the data range nearing a mean value, while employing a low resolution, i.e., fewer quantized values, for sparsely distributed data range. The densely distributed data range contains more weights, and adopting high resolution can significantly reduce the quantization errors for quantizing these weights. Correspondingly, employing low resolution to the sparsely distributed data range pays almost no impact on the quantization errors, because there are few weights in this range. Namely, fitting distribution can reduce quantization errors by finding a set of quantized values which distribute close to the weight distribution. Deviating weight distribution can be harmful to quantization errors, thereby decreasing model accuracy. Therefore, a desirable quantization should meet two requirements: the difference between the two distributions of quantized values and weights, and the

- Cheng Gong, Kunpeng Xie, and Zongming Jin are with the College of Computer Science, Nankai University, Tianjin 300350, China, and also with the Tianjin Key Laboratory of Network and Data Security Technology, Tianjin 300350, China. E-mail: {cheng-gong, xkp, zongming_jin}@mail.nankai.edu.cn.
- Ye Lu and Tao Li are with the College of Computer Science, Nankai University, Tianjin 300350, China, and also with the Tianjin Key Laboratory of Network and Data Security Technology, State Key Laboratory of Computer Architecture (ICT,CAS), Tianjin 300350, China. E-mail: {luye, litao}@nankai.edu.cn.
- Yanzhi Wang is with the Department of Electrical and Computer Engineering, Khoury College of Computer Science (Affiliated), Northeastern University, Boston, MA 02115 USA. E-mail: yanzhi.wang@northeastern.edu.

Manuscript received 30 Aug. 2021; revised 29 Oct. 2021; accepted 12 Nov. 2021. Date of publication 22 Nov. 2021; date of current version 23 May 2022.

This work was supported in part by the National Key Research and Development Program of China under Grant 2018YFB2100300, the National Natural Science Foundation under Grant 62002175, the Natural Science Foundation of Tianjin under Grant 19JCQNJC00600, the State Key Laboratory of Computer Architecture (ICT,CAS) under Grants CARCHB202016, CARCH201905, and the Innovation Fund of Chinese Universities Industry-University-Research under Grant 2020HYA01003.

(Corresponding author: Tao Li.)

Recommended for acceptance by A.J. Peña, M. Si and J. Zhai.

Digital Object Identifier no. 10.1109/TPDS.2021.3129615

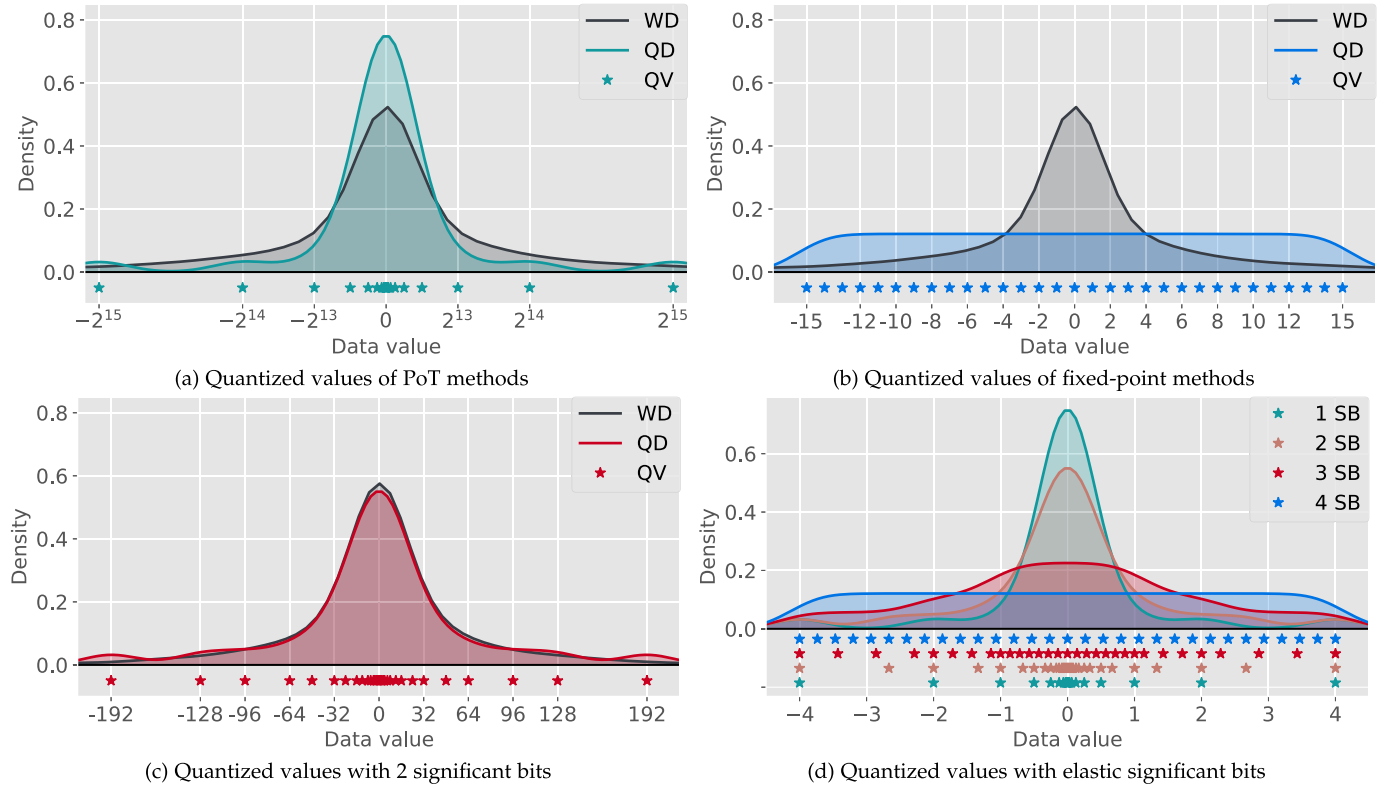


Fig. 1. Distributions of different quantized values. *WD* is the weight distribution of the first convolutional layer of ResNet18 [1] with floating-point data type. *QD* denotes quantized value distribution and *QV* indicates quantized values. *SB* indicates significant bits.

number of bits used should both be as small as possible. In other words, to maintain the original features of DNNs, quantized values should fit the distribution of the original values by using fewer bits. The closer the two distributions, the higher the quantized DNN accuracy.

Most of the previous studies cannot fit the distribution of the original weights/activation well. As shown in Fig. 1, the distribution of original weights is typically bell-shaped [16] in the first convolutional layer of ResNet18 [1]. With five bits for quantization, the distribution of the quantized values of the power of two (PoT) scheme is sharply long-tailed [13], [17], [18], and that of the fixed-point method is uniform [7], [19], [20], as shown in Figs. 1a and 1b, respectively. We can observe that both PoT and the fixed-point deviate from the bell-shaped distribution. Although recent studies [14], [15], [21] can reduce accuracy degradation by superposing a series of binaries/PoTs, they still encounter limitations in fitting the distributions of weights/activation well.

We observe that the number of significant bits among the given can impact the distribution of quantized values. Significant bits in computer arithmetic refer to the bits that contribute to the measurement resolution of the values. Values represented by different numbers of significant bits and offsets can present different distributions. It implies that we can fit the distributions of weights and activation with elastic significant bits to reduce the accuracy degradation. In addition, fewer significant bits of operands incur less multiplication complexity, because the multiply operation is implemented by a group of shift-accumulate operations. The group size is proportional to the number of significant bits of operands. Therefore, the number of significant bits

representing weights/activation can directly affect the accuracy and efficiency of the quantized DNN inference. This motivates us to utilize the significant bits as a key technique to fit the distribution.

To demonstrate the advantages of significant bits, we use an example to further clarify our motivation. Specifying two significant bits within the five given bits for quantization, we first construct two value sets in which the number of significant bits of values is no more than two: $\{0,1\}, \{2,3\}$. Then, we enumerate the shifted sets of $\{2,3\}$, including $2 \times \{2,3\}$, $2^2 \times \{2,3\}$, \dots , $2^6 \times \{2,3\}$, until the number of elements of all sets reaches 2^4 . After merging these sets and adding the symmetric negative values, we can obtain the set of $\{0, \dots, \pm 32, \pm 48, \pm 64, \pm 96, \pm 128, \pm 192\}$ with $2^5 - 1$ elements. As shown in Fig. 1c, the distribution of the constructed set can fit the weight distribution well, which implies that quantizing DNNs using these quantized values can yield higher accuracy. In addition, multiplication of the values from the set only requires two shift-accumulate operations, as shown in Fig. 2, which is the same as the 2-bit quaternary multiplication [8], [9]. Namely, quantization using two significant bits can achieve high accuracy as the 5-bit

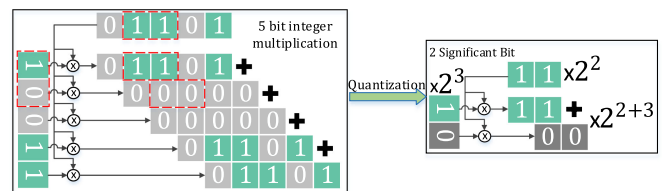


Fig. 2. Quantizing values with elastic significant bits can reduce multiplication consumption.

fixed-point scheme, but only requires a similar consumption with the multiplication of 2-bit quaternary.

In this study, we propose elastic significant bit (ESB) quantization to achieve high accuracy with few resources by controlling the number of significant bits of quantized values. ESB involves three parts: 1) the mathematical formula describing values with a flexible number of significant bits unifies various quantized value set featuring diverse distributions, as shown in Fig. 1d; 2) the distribution difference aligner (DDA) evaluating the matching degree between distributions seeks the desirable quantized value set which fits the original distribution; 3) ESB float format coding quantized values in hardware implementation realizes an efficient multiplication computation with low resource consumption. Our contributions are summarized as follows:

- To the best of our knowledge, we propose the first unified mathematical expression to describe the quantized values with a flexible number of significant bits, and design a hardware-friendly projection for mapping from full-precision values to these quantized values at runtime. This formula unifies diverse quantized values, including both PoT and fixed-point forms, and make the values suitable for various distributions, such as the uniform and the bell-shaped ones.
- We introduce DDA to evaluate the matching degree of distributions before and after quantization. Based on DDA, the desirable distribution of quantized values, fitting the bell-shaped weights and activation of DNNs, can be found offline without incurring additional computations.
- We design ESB float format to implement the accelerator based on FPGAs, which codes and stores the fraction and exponent of ESB quantized values separately. This format enables the accelerator to employ low-precision fraction multiplication and shift operation, thus reducing resource consumption and improving energy efficiency.
- Experimental evaluations on the ImageNet dataset with AlexNet, ResNet18, and MobileNetV2 demonstrated that ESB achieves higher accuracy than 15 state-of-the-art approaches by 4.78%, 1.92%, and 3.56% on average, respectively. The ESB accelerator on FPGAs can achieve a peak performance of 10.95 GOPS/kLUTs without DSPs and improve the energy efficiency by up to 65 \times , 11 \times and 26 \times compared with CPU, GPU, and state-of-the-art accelerators, respectively.

2 RELATED WORKS

DNNs have been widely applied in various areas [22], [23], [24], [25], but the high-bitwidth floating-point operations in DNNs affect the inference efficiency seriously. Early studies on quantization design ultra-low-bitwidth methods to improve the inference efficiency of DNNs, such as binary and ternary methods [4], [26]. These studies achieve high computing efficiency with the help of a low bitwidth but failed to guarantee accuracy. Thus, researchers have focused on multi-bit methods to prevent accuracy degradation.

Linear Quantization. The linear quantization quantizes the data into consecutive fixed-point values with uniform intervals. Ternarized hardware deep learning accelerator (T-DLA) [5] argues that there are invalid bits in the heads of the binary strings of tensors. It drops both head invalid bits and tail bits in the binary strings of activation. Two-step quantization (TSQ) [26] uniformly quantizes the activation after zeroing the small values and then quantizes the convolutional kernels into the ternaries. Blended coarse gradient descent (BCGD) [20] uniformly quantizes all weights and the activation of DNNs into fixed-point values with shared scaling factors. Parameterized clipping activation (PACT) [27] parameterizes the quantizing range. Furthermore, it learns the proper range during DNN training and linearly quantizes activation within the learned range. Quantization-interval-learning (QIL) [19] parameterizes the width of the intervals of linear quantization and learns the quantization strategy by optimizing the DNN loss. Differentiable Soft Quantization (DSQ) [7] proposes a differentiable linear quantization method to resolve the non-differentiable problem, applying the *tanh* function to approximate the step function of quantization by iterations. Hardware-Aware Automated Quantization (HAQ) [28] leverages reinforcement learning (RL) to search the desired bitwidths for different layers and then linearly quantizes the weights and activation of the layers into fixed-point values.

Despite these significant advances in linear quantization, it is difficult to achieve the desired trade-off between efficiency and accuracy owing to the ordinary efficiency of fixed-point multiplication and mismatching nonuniform distributions.

Non-Linear Quantization. Non-linear quantization projects full-precision values into low-precision values that have non-uniform intervals. Residual-based methods [1], [29] iteratively quantize the residual errors produced by the last quantizing process to binaries with full-precision scaling factors. Similar methods, such as learned quantization networks (LQ-Nets) [30], accurate-binary-convolutional network (ABC-Net) [21], and AutoQ [31], quantize the weights or activation into the sum of multiple groups of binary quantization results. The multiply operations of the quantized values of residual-based methods can be implemented as multiple groups of bit operations [4] and their floating-point accumulate operations. The computing bottleneck of their multiply operations depends on the number of accumulate operations, which is proportional to the number of quantizing bits. Float-based methods refer to the separate quantization of the exponent and fraction of floating-point values. The block floating-point (BFP) [32] and Flexpoint [33] retain the fraction of all elements in a tensor/block and share the exponent to reduce the average number of quantizing bits. Posit schemes [34], [35] divide the bit string of values into four regions to produce a wider dynamic range of weights: sign, regime, exponent, and fraction. Adaptive-Float [36] quantizes full-precision floating-point values into a low-precision floating-point format representation by discretizing the exponent and fraction parts into low-bitwidth values, respectively. Float-based methods usually present unpredictable accuracy owing to a heuristical quantizing strategy and zero value absence in the floating-point format.

Few studies have investigated reducing the accuracy degradation of float-based methods.

Quantization Fitting Distribution. An efficient way to reduce accuracy degradation by fitting the distributions in multi-bit quantization has been extensively investigated. Studies on weight sharing, such studies focused on K-means [16], [37], take finite full-precision numbers to quantize weights, and store the frequently used weights. Although the shared weights can reduce the storage consumption, it is difficult to reduce the multiplication complexity of full-precision quantized values. Recent studies [8], [9] attempt to design a better linear quantizer by adjusting the step size of the linear quantizer to match various weight distributions, thus reducing accuracy degradation. However, the linear quantizer cannot fit the nonuniform weights well [15], which inevitably degrades the accuracy. Quantizing the weights into PoT is expected to fit the non-uniform distributions while eliminating bulky digital multipliers [13]. However, PoT quantization cannot fit the bell-shaped weights well in most DNNs [15], as shown in Fig. 1a, thus resulting in significant accuracy degradation [17], [18]. AddNet [14] and additive powers-of-two quantization (APoT) [15] exhaustively enumerate the combinations of several computationally efficient coefficients, such as reconfigurable constant coefficient multipliers (RCCMs) and PoTs, to fit the weight distribution while maintaining efficiency. However, exhaustive enumeration has a high complexity of $O(n!)$. In addition, they require expensive look-up operations to map the full-precision weights into their nearest quantized values, that is, the universal coefficient combinations. These defects can significantly delay model training.

3 ELASTIC SIGNIFICANT BIT QUANTIZATION

In this section, we introduce the ESB to explain how the quantized values can be controlled at the bit-wise level to fit various distributions. First, we introduce some preliminaries and then expound our ESB from two aspects: quantized value set and cheap projection operator. Second, we introduce the DDA to fit the original value distribution and reduce quantization errors. Finally, We present how ESB can be exploited to train DNNs.

3.1 Preliminaries

Notations. For simplicity, we use weight quantization as an example to describe the quantization process. We denote the original weights of a layer in DNNs as $W_f \in \mathbb{R}^d$, and d indicates its dimension. W_q represents the quantized weights.

Quantized Value Set and Projection Operator. A general quantization process projects full precision values into quantized values through the projection operator Π , and these quantized values comprise set Q . Π impacts the projection efficiency, and Q reflects the distribution law of quantized values. The quantization can be defined as follows:

$$W_q = \Pi_Q(W_f). \quad (1)$$

Here, Π projects each original value in W_f into a quantized value in Q and stores it in W_q . For *linear quantization*, all weights are quantized into fixed-point values as follows:

$$Q_l(\alpha, b) = \{0, \pm\alpha, \pm2\alpha, \pm3\alpha, \dots, \pm(2^{b-1} - 1)\alpha\}. \quad (2)$$

For PoT, all weights are quantized into values in the form of PoT as follows:

$$Q_p(\alpha, b) = \{0, \pm2^0\alpha, \pm2^1\alpha, \pm2^2\alpha, \dots, \pm2^{2^{b-1}-1}\alpha\}. \quad (3)$$

Here, α is the scaling factor, and b represents the number of bits for quantization. $Q_l(\alpha, b)$ and $Q_p(\alpha, b)$ cannot fit the bell-shaped distribution well, as shown in Figs. 1b and 1a, leading to inevitable accuracy degradation. Projection Π adopts the nearest-neighbor strategy. It projects each value in W_f into the nearest quantized value in Q measured by the euclidean distance. Because $Q_l(\alpha, b)$ consists of fixed-point values with uniform intervals, linear quantization projection is implemented economically by combining truncating with rounding operations [6], [28]. PoT projection can be implemented using a look-up function [15], [17], [18]. However, the look-up function has a time complexity of $O(n)$. It involves $d * (2^b - 1)$ comparison operations for projecting W_f into $Q_p(\alpha, b)$. Hence, PoT projection is more expensive than linear quantization projection.

3.2 Quantized Value Set

First, we define the quantized value set and the projection operator for the ESB as follows.

$$\begin{aligned} W_q &= \Pi_{Q_e(\alpha, b, k)}(W_f) \\ \text{s.t. } Q_e(\alpha, b, k) &= \alpha \times (\cup_{i=0}^N \xi_i \times \Omega_i), \quad N = 2^{b-k-1} - 1. \end{aligned} \quad (4)$$

Here, $Q_e(\alpha, b, k)$ is the quantized value set, and α is a scaling factor. Ω_i ($i = 0, 1, \dots, N$) is the set of values with the number of significant bits no more than $k + 1$, and ξ_i is a scalar. We call Ω_i the fraction set and ξ_i the shift factor, and we define them as follows.

$$\begin{aligned} \Omega_i &= \begin{cases} \{0, 1, 2, \dots, 2^k - 1\}, & i = 0 \\ \{2^k, 2^k + 1, \dots, 2^k + (2^k - 1)\}, & i > 0 \end{cases} \\ \xi_i &= \begin{cases} 2^{-k}, & i = 0 \\ 2^{i-k-1}, & i > 0 \end{cases} \end{aligned} \quad (5)$$

Here, k is an integer parameter for controlling the number of significant bits of quantized values. b is the number of bits for quantization as mentioned previously. Obviously, the number of significant bits of the values in Ω_i is no more than $k + 1$ and ξ_i is a PoT value. $Q_e(\alpha, b, k)$ is constructed using the shifted fraction sets $\xi_i \times \Omega_i$. For ease of notation, we used $\pi = 2^k \times \{0, \xi_1, \xi_2, \dots, \xi_N\}$ to denote the PoT values in $Q_e(\alpha, b, k)$. We present the visualized format of the values in $Q_e(\alpha, b, k)$ in Fig. 3.

We would like to point out that the above mathematical formulas unify all types of distributions, and this is one of our important contributions. In particular, linear quantization can be regarded as a typical case of ESB quantization when the condition $k = b - 2$ is satisfied. Similarly, PoT is equal to the ESB when $k = 0$. Ternary is a special case of ESB when $k = 0$ and $b = 2$ are both satisfied.

The multiplication of elements in $Q_e(\alpha, b, k)$ can be converted to the corresponding multiplication of elements in Ω_i and shift operation (multiply shift factor ξ_i). Therefore, the

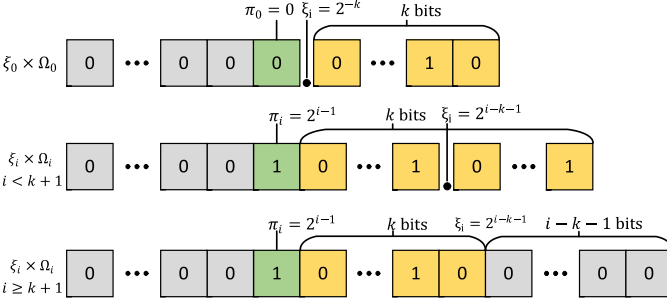


Fig. 3. Format of quantized values in ESB.

overhead of matrix multiplications in quantized DNNs with ESB relies on the number of significant bits of values in Ω_i . The number is controlled with a configurable parameter k instead of the given bits b . In other words, controlling parameter k in the ESB can directly impact the computing efficiency of the discretized DNNs. In addition, we observe that the elastic k can facilitate ESB to be suitable for various distributions, ranging from the sharply long-tailed ($k = 0$) to the uniform ($k = b - 2$) one, including the typical bell-shaped distribution, as shown in Fig. 1d.

Consequently, we can conclude that elastic k has two major effects: fitting the distribution and improving computing efficiency. By controlling the number of significant bits, the ESB can achieve a trade-off between accuracy and efficiency.

3.3 Projection Operator

Next, we shift our attention to projection operator Π . We define the full precision value v for the projection operator. v enables a look-up function $P(\cdot)$ based on the nearest neighboring strategy as follows:

$$P(v) = \begin{cases} \pi_i, & v \in [\pi_i - \frac{\xi_i-1}{2}, \pi_i + \frac{\xi_i}{2}), i = 0, 1, \dots, N \\ \tau_i, & v \in [\tau_i - \frac{\xi_i}{2}, \tau_i + \frac{\xi_i}{2}), i = 0, 1, \dots, N \end{cases} \quad (6)$$

$s.t. v \in [-C, C]$

Here, $\tau_i \in \{x : x \in \xi_i \times \Omega_i, x \neq \pi_i\}$, $C = \max(\xi_N \times \Omega_N)$, and $\max(\cdot)$ returns the maximum element. The look-up function

in (6) is convenient for formula expression but expensive in practice, as mentioned in Section 3.1.

Inspired by the significant bits, we observe that the essence of ESB projection is to reserve limited significant bits and round the high-precision fraction of values. Therefore, we combine the *shift* operation with the *rounding* operation to implement $P(\cdot)$ and reduce computational overhead. Specifically, for a value $v \in [-C, C]$, we seek its most significant 1 bit: $n = \max(\{i : v \geq 2^i\})$. Let $x(i)$ represent the i th binary valued in $\{0, 1\}$, v can be divided into two parts. One is $\sum_{i=n-k}^n 2^i \cdot x(i)$, which has $k+1$ significant bits and should be reserved. The other part is $\sum_{i=-\infty}^{n-k-1} 2^i \cdot x(i)$, which represents the high-precision fraction and should be rounded. Therefore, $P(\cdot)$ can be refined as

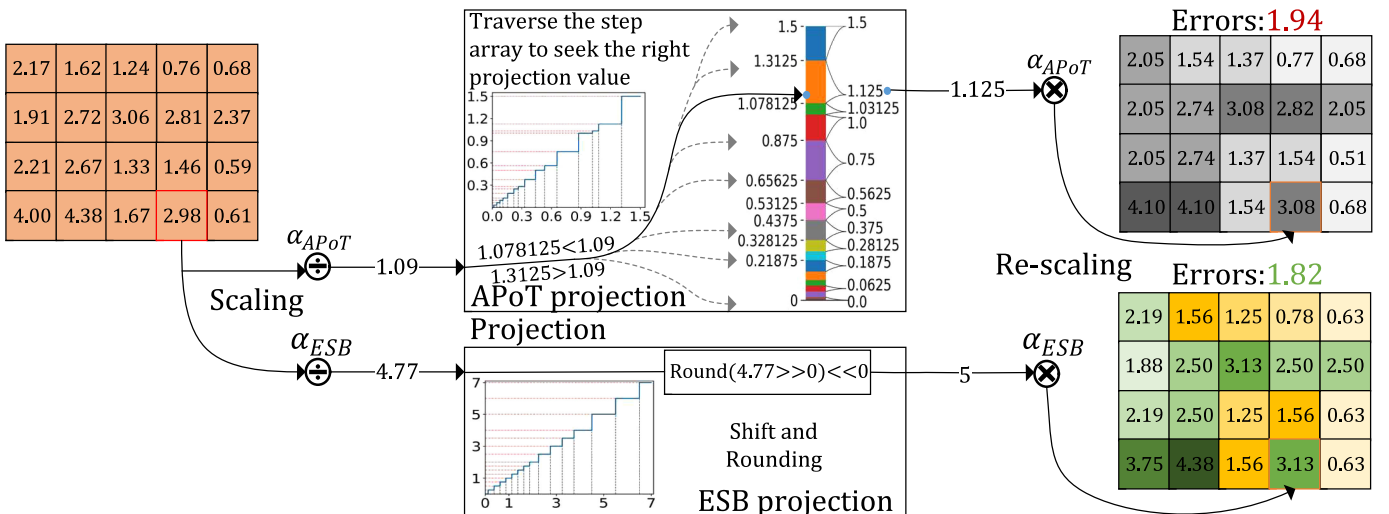
$$P(v) = R(v \gg (n-k)) \ll (n-k) \quad s.t. v \in [-C, C]. \quad (7)$$

Here, the symbols \ll and \gg represent the left-shift operation and right-shift operation, respectively. $R(\cdot)$ represents *rounding* operation. All of them are economical computations because they require only one clock cycle computing in modern CPU architectures. Based on (7), the ESB projection operator Π can be redesigned as follows:

$$\Pi_{Q_e(\alpha, b, k)}(W_f) = \{\alpha P(v) : v \in [W_f/\alpha]_{-C}^C\}. \quad (8)$$

Here, $[\cdot]_{-C}^C$ truncates the data in the range of $[-C, C]$. Consequently, this projection operator is beneficial for quantizing activation and weights. In particular, the multiplications recognized as the performance bottleneck can also be processed directly, without a projection delay.

To demonstrate the superiority of our method theoretically, we provide a comparison figure to show the projection process of APoT and ESB. As shown in Fig. 4, given a vector consists of 20 elements and 5 bits for quantization, APoT (with 2 groups) requires $s = 2^5 - 1$ memory space to pre-store the step-values of the look-up function. For the input value 2.98, APoT quantizes it into 3.08 through three stages: scaling 2.98 to 1.09 with α_{APoT} , costing $s/2$ times on average to look up the right projection value 1.125, and re-scaling 1.125 to 3.08. Finally, APoT quantizes the original

Fig. 4. Projection of APoT [15] (2 groups with scaling factor α_{APoT}) and ESB ($k = 2$ and α_{ESB}) with five given bits.

vector to a quantized vector with an error of 1.94, and the time and space complexities of APoT are both $O(s)$. In contrast, ESB ($k = 2$) does not need to store step-values in advance. ESB quantizes the value 2.98 through three stages as well: scaling 2.98 to 4.77 with α_{ESB} , obtaining the projection value 5 through cheap shift and rounding operations, and re-scaling 5 to 3.13. ESB finally produces a quantized vector with an error of 1.82, and its time and space complexities are only $O(1)$. Thus, we can conclude that the APoT operator implies $O(s)$ (for storing s steps) for time and space complexity. This leads to high computing overheads compared with our ESB.

3.4 Distribution Difference Aligner

Assuming that W_f is sampled from a random variable t . $p(t; \theta_{W_f})$ is the probability of t , and θ_{W_f} is the parameter vector. For simplicity, we divide this distribution into two parts: the negative ($t < 0$) and the non-negative ($t \geq 0$), because the quantized value set is usually symmetric. Next, we define the difference $D(\alpha, b, k)$ between the two distributions on the non-negative part as follows:

$$D(\alpha, b, k) = \sum_{q \in Q_e^*(\alpha, b, k)} \int_{t \in \mathcal{S}} (t - q)^2 p(t; \theta_{W_f}) dt$$

$$s.t. \mathcal{S} = \begin{cases} [0, \alpha \frac{\xi_0}{2}), & q = 0 \\ [q - \alpha \frac{\xi_N}{2}, \infty), & q = C \\ [q - \alpha \frac{\xi_{i-1}}{2}, q + \alpha \frac{\xi_i}{2}], & q = \alpha \pi_i, i = 0, 1, \dots, N \\ [q - \alpha \frac{\xi_i}{2}, q + \alpha \frac{\xi_i}{2}), & q = \alpha \tau_i, i = 0, 1, \dots, N \end{cases} \quad (9)$$

The symbol $*$ indicates the non-negative part of the set. Generally, for a symmetric distribution, we can utilize $D(\alpha, b, k)$ above instead of the entire distribution to evaluate distribution difference. The challenge of evaluating and optimizing $D(\alpha, b, k)$ is to estimate the probability distribution $p(t; \theta_{W_f})$ of weights. Here, we employ data normalization (DN) and parameterized probability density estimation to address this challenge.

Normalization. Data normalization (DN) has been widely used in previous works [15], [38], refining data distribution with zero mean and unit variance. Let μ and σ denote the data mean and standard derivation, respectively, weight normalization can be defined as follows:

$$\hat{W}_f = \frac{W_f - \mu_{W_f}}{\sigma_{W_f} + \epsilon}$$

$$s.t. \begin{cases} \mu_{W_f} = \frac{1}{d} \sum_{j=1}^d W_f(j) \\ \sigma_{W_f} = \sqrt{\frac{1}{d} \sum_{j=1}^d (W_f(j) - \mu_{W_f})^2} \end{cases} \quad (10)$$

Here, $\epsilon = 1e - 7$ is a small value to avoid overflow of division. Estimating the mean and standard deviation of activation results in heavy workloads. We employ an *exponential moving average* [38] for parameter estimation of activation on the entire training set in the model training phase. For the i th training batch with activation, A_{fi} , $\mu_{A_{fi}}$, and $\sigma_{A_{fi}}$ are computed as follows.

$$\begin{cases} \mu_{A_f} = (1 - \gamma)\mu_{A_f} + \gamma\mu_{A_{fi}} \\ s.t. \mu_{A_{fi}} = \frac{1}{d} \sum_{j=1}^d A_{fi}(j) \\ \sigma_{A_f} = (1 - \gamma)\sigma_{A_f} + \gamma\sigma_{A_{fi}} \\ s.t. \sigma_{A_{fi}} = \sqrt{\frac{1}{d} \sum_{j=1}^d (A_{fi}(j) - \mu_{A_{fi}})^2} \end{cases} \quad (11)$$

Here, γ is a momentum coefficient and is typically set to 0.9 or 0.99. μ_{A_f} and σ_{A_f} are the estimated mean and standard deviation of activation on the entire training set, respectively. The definitions of activation normalization for A_{fi} in the training and test phases are as follows:

$$\text{Training: } \hat{A}_{fi} = \frac{A_{fi} - \mu_{A_{fi}}}{\sigma_{A_{fi}} + \epsilon} \quad \text{Test: } \hat{A}_{fi} = \frac{A_{fi} - \mu_{A_f}}{\sigma_{A_f} + \epsilon}. \quad (12)$$

Algorithm 1. DNN Training With ESB

Input: Iteration times T . Dataset D . DNN depth L . Learning rate η .

Output: The trained DNN model.

```

1: for t = 1 to T do
2:    $A_f(1) = D(t)$ 
3:   for i = 1 to L do
4:      $\hat{A}_f(i) = \text{DN}(A_f(i))$ ,  $\hat{W}_f(i) = \text{DN}(W_f(i))$ 
5:      $A_q(i) = \text{ESB}(\hat{A}_f(i))$ ,  $W_q(i) = \text{ESB}(\hat{W}_f(i))$ 
6:      $A_f(i+1) = \text{Forward}(A_q(i), W_q(i))$ 
7:   end for
8:   Computing the model loss  $Loss$ .
9:    $G_A(L+1) = \frac{\partial Loss}{\partial A_f(L+1)}$ 
10:  for i = L to 1 do
11:     $G_A(i) = G_A(i+1) \cdot \frac{\partial A_f(i+1)}{\partial A_q(i)} \cdot \frac{\partial \hat{A}_f(i)}{\partial A_f(i)}$ 
12:     $G_W(i) = G_A(i+1) \cdot \frac{\partial A_f(i+1)}{\partial W_q(i)} \cdot \frac{\partial \hat{W}_f(i)}{\partial W_f(i)}$ 
13:  end for
14:  for l = 1 to L do
15:     $W_f(i) = W_f(i) - \eta \cdot G_W(i)$ 
16:  end for
17: end for
```

Resolving extremum. Without loss of generality, we assume that all weights and activation of each layer in the DNNs satisfy the normal distribution [8], [26]. Through the above normalization steps, the distribution parameter vector θ_{W_f} is $[0, 1]^T$. Thus, the probability distribution is: $p(t; \theta_{W_f}) = \mathcal{N}(t; 0, 1)$. After defining $p(t; \theta_{W_f})$, all the parts of (9) are differentiable, and our optimization can be defined as follows:

$$\alpha_{b,k}^* = \arg \min_{\alpha_{b,k}} (D(\alpha, b, k)), \quad (13)$$

As shown in Fig. 5, we draw the $D(\alpha, b, k)$ curves for α across various conditions. The horizontal axis represents the value of α , and the vertical axis represents $D(\alpha, b, k)$. When α values are obtained from the given range $(0, 3]$, there exist some convex curves such as $D(\alpha, 2, 0)$, $D(\alpha, 3, 0)$, $D(\alpha, 3, 1)$, $D(\alpha, 4, 1)$, etc. These convex curves imply that we can resolve their extremums through mathematical methods as optimal solutions. For non-convex curves, we can still obtain their extremums as feasible solutions which are

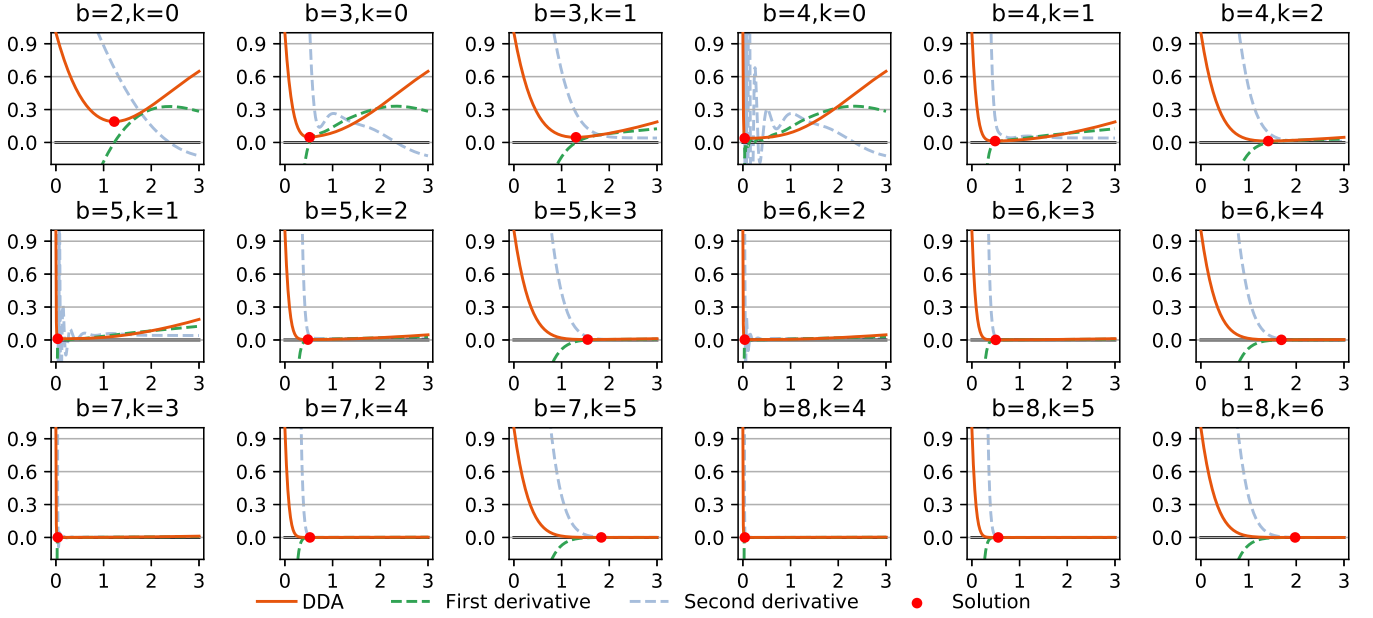


Fig. 5. $D(\alpha, b, k)$ with different α and its first and second derivative curves for different values of b and k .

close to the optimal solutions. Therefore, we can directly resolve the extremum of $D(\alpha, b, k)$ to obtain a solution $\alpha_{b,k}^*$ to achieve accurate quantization. We have computed all desirable solutions $\alpha_{b,k}^*$ for different values of b and k . The solutions are presented in Table 1.

3.5 Training Flow

For DNN training, we present the overall training flow as shown in Fig. 6. Let DNN depth be L , and the 32-bit floating-point weights and activation of the i th layer be $W_f(i)$ and $A_f(i)$, respectively. We embed the ESB quantization before the multiplier accumulator (MAC) and quantize $W_f(i)$ and $A_f(i)$ into low-precision $W_q(i)$ and $A_q(i)$, respectively. Then, $W_q(i)$ and $A_q(i)$ are fed into the MAC for acceleration. In backward propagation, we employ straight-through-estimation (STE) [6] to propagate the gradients. The details are described in Algorithm 1.

4 HARDWARE IMPLEMENTATION

The efficient multiplication implementation of ESB quantized values relies on a flexible MAC design, such as low-bit-width multiplication and shift operation designs. However, ESB quantized values cannot be processed efficiently on current general-purpose CPU/GPU platforms, since they do not provide the corresponding instructions for processing these values with flexible bitwidth configurations.

Consequently, in this section, we introduce the hardware architecture design of the ESB based on the Field Programmable Gate Array (FPGA) platform. We first present the float format for representing ESB quantized values and explain the multiplication process with ESB float format values. Then, we present the overall hardware architecture design, which includes 1) the general computing logic of DNNs (taking AlexNet as an example) as shown in Fig. 9, 2) hardware architecture as shown in Fig. 10, and 3) timing graph as shown in Fig. 11.

4.1 ESB Float Format

The ESB float format was designed as shown in Fig. 7. It is similar to AdaptiveFloat [36] (IEEE 754 format), that is, 1 bit for the sign, k bits for the fraction, and $b - k - 1$ bits for the exponent. Let e be the exponent value, f be the fraction value, and sb be the sign bit (0 or 1); thus, a quantized value q can be represented as follows:

$$q = \begin{cases} (-1)^{sb} * 0.f, & e = 2^{b-k-1} - 1 \\ (-1)^{sb} * 2^e * 1.f, & \text{otherwise} \end{cases} \quad (14)$$

Here, we use the representations with a maximum exponent value of $2^{b-k-1} - 1$ to represent small numbers around zero. The difference from AdaptiveFloat is that the design of the ESB float format presents smaller and gradually progressive increasing gaps around zero. AdaptiveFloat cannot represent zero accurately, or it forces the absolute minimum values to zero, resulting in evident larger gaps around zero and losing large information, as shown in Fig. 8. By contrast, ESB can present progressively increasing gaps around zero, thus facilitating a smooth distribution around zero and better fitting the weight/activation distributions. The multiplication of two given values, such as q_i with sb_i, e_i, f_i and q_j with sb_j, e_j, f_j , can be derived as follows:

$$q_i \cdot q_j = (-1)^{sb_i \oplus sb_j} \cdot (f < e) \quad (15)$$

$$s.t. \begin{cases} f = (\zeta_i + 0.f_i) \cdot (\zeta_j + 0.f_j) \\ e = e_i \zeta_i + e_j \zeta_j \\ \zeta_i = \mathbb{I}_{e_i \neq 2^{b-k-1}-1}, \zeta_j = \mathbb{I}_{e_j \neq 2^{b-k-1}-1} \end{cases}$$

Here, $\zeta_i + 0.f_i$ and $\zeta_j + 0.f_j$ are the fraction values. Therefore, the b -bit multiplication of the ESB float format can be implemented by a $k + 1$ bit multiply operation at most. $k + 1$ is the number of significant bits of the fraction values.

4.2 Computing Logic

We use AlexNet as an example to introduce the typical computing logic of a quantized DNN with ESB. AlexNet has five

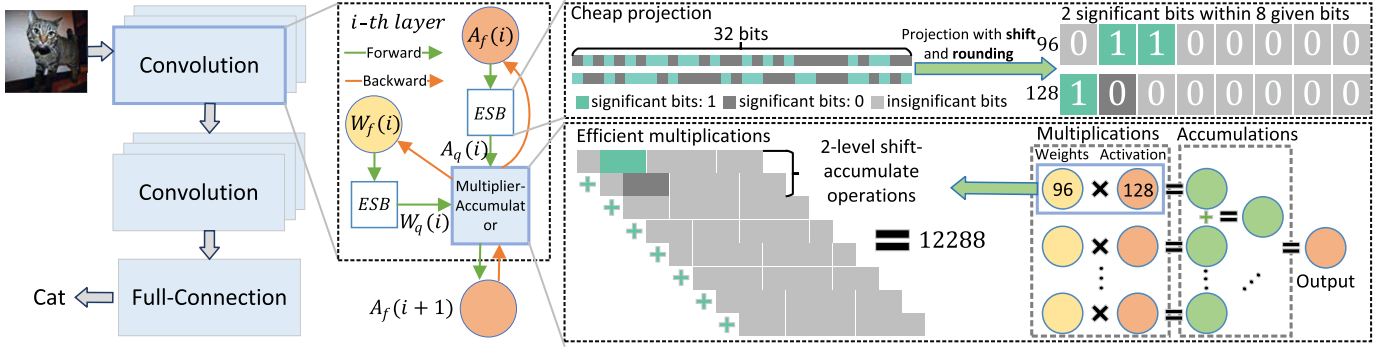


Fig. 6. Overall training flow of DNNs with ESB quantization.

convolutional layers and three fully connected layers. A convolutional layer usually consists of four stages: convolution, pooling (max-pooling, MP), activation (ReLU), and batch normalization (BN). In ESB, the scale stage (Scale), data normalization (DN) (Section 3.4) stage, and ESB quantization (in (8)) stage are attached to quantize the activation values. Therefore, the l th convolution layer of AlexNet includes seven stages in our implementation, as shown in Fig. 9.

In the convolution stage, we employ the weight $W_q(l) \in Q_e(\alpha_{W_q(l)}, b, k)^{K \times K \times c_{in} \times c_{out}}$ to filter the input feature map $A_q(l) \in Q_e(\alpha_{A_q(l)}, b, k)^{h \times w \times c_{in}}$ and output $A_f(l+1) \in R^{h \times w \times c_{out}}$. $W_q(l)$ and $A_q(l)$ are the ESB float-format tensors. In particular, a convolution operation traverses the input feature map $A_q(l)$ and multiplies accumulating c_{in} channel features and the corresponding c_{in} kernels to produce the output feature map $A_f(l+1)$. The total number of operations for a convolution computation is up to $h \cdot w \cdot K \cdot K \cdot c_{in} \cdot c_{out}$. The following MP operation employs a window of size $p \times p$ to slide on the feature map by the stride step of s and computes the maximum value within the window during sliding. If s is not divisible by $h - p$ or $w - p$, the sliding window cannot be aligned with the feature map. We captured the intersection window to compute the MP result. After the MP operation, the size of the feature map $A_f(l+1)$ degrades to $(\lceil \frac{h}{s} \rceil, \lceil \frac{w}{s} \rceil, c_{out})$. The following five stages only involve injective operations and are defined as shown in Fig. 9. These stages can be fused in the model inference.

4.3 Hardware Architecture

Our accelerator architecture is shown in Fig. 10, which has two modules: convolution and post-processing. The convolution module processes the computation-intensive convolution operation, whereas the post-processing module processes all the remaining stages, including MP, ReLU, Scale, BN, DN, and ESB quantization.

In the convolution module, we split the original feature map and output feature map into small tensors with sizes of

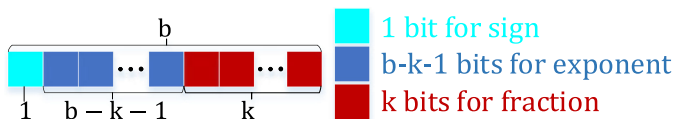


Fig. 7. ESB float format.

(T_h, T_w, T_n) and (T_h, T_w, T_m) as shown in Fig. 9, and split the computing operations into $\lceil \frac{h}{s} \rceil \cdot \lceil \frac{w}{s} \rceil \cdot \lceil \frac{c_{in}}{T_n} \rceil \cdot \lceil \frac{c_{out}}{T_m} \rceil$ groups of small convolution. To hide the data transmission time, we employ dual ping-pong input buffers (0/1) of size (T_h, T_w, T_n) and dual ping-pong output buffers (0/1) of size (T_h, T_w, T_m) for a small convolution computation, as shown in Fig. 10. Then, we design $T_m \cdot T_n$ MAC units, which can execute $2 \cdot T_m \cdot T_n$ multiply-accumulating operations in parallel. Therefore, we can reduce the computing time of the convolution stage by $T_m \cdot T_n$ times and accelerate the convolution inference.

In the post-processing module, we compress the six stages into three steps by fusing some stages into one, as shown in Fig. 10. The "MP&ReLU" step fuses MP and ReLU operations. MP takes a sliding window with a pooling size of $p \times p$ and a stride step of s to process the small output tensor of size (T_h, T_w, T_m) . If the sliding window of the MP cannot be aligned with the tensor, we cache their intersection window and delay the pooling operation until the next computing trip. The following ReLU eliminates negative MP outputs. Then, we fuse the Scale, BN, DN stages, and the division part of ESB quantization into a single step of "BN and DN" as follows:

$$\text{DN}(\text{BN}(\text{Scale}(x))) / \alpha_{A_q}(l+1) = ax + b$$

$$s.t. \begin{cases} a = \frac{\alpha_{W_q(l)} \cdot \alpha_{A_q(l)} \cdot \lambda_{BN}}{\alpha_{A_q(l+1)} \cdot \sigma_{A_f(l+1)} \cdot \sigma_{BN}} \\ b = \frac{\sigma_{BN} \cdot \beta_{BN} - \mu_{BN} \cdot \lambda_{BN} - \sigma_{BN} \cdot \mu_{A_f(l+1)}}{\alpha_{A_q(l+1)} \cdot \sigma_{A_f(l+1)} \cdot \sigma_{BN}} \end{cases} \quad (16)$$

Here, $\alpha_{A_q(l)}$, $\alpha_{W_q(l)}$ are the scaling factors of $A_q(l)$ and $W_q(l)$ of the l th layer, respectively. $\alpha_{A_q(l+1)}$ is the scaling factor for $A_q(l+1)$, which is the input of the $(l+1)$ -th layer. μ_{BN} , σ_{BN} , λ_{BN} , and β_{BN} are the parameters of BN, $\mu_{A_f(l+1)}$ and $\sigma_{A_f(l+1)}$ are the parameters of DN. Since all of these parameters are constant in the model inference, we compute a and b offline

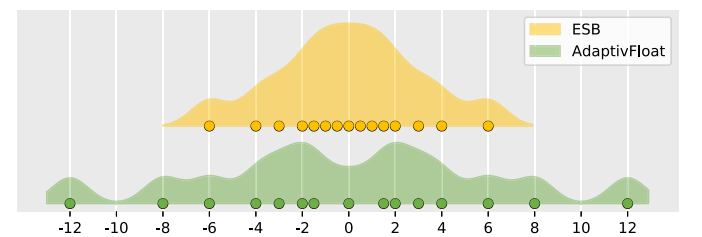


Fig. 8. Quantized values of ESB and AdaptivFloat with the same 1-bit sign, 2-bit exponent and 1-bit fraction.

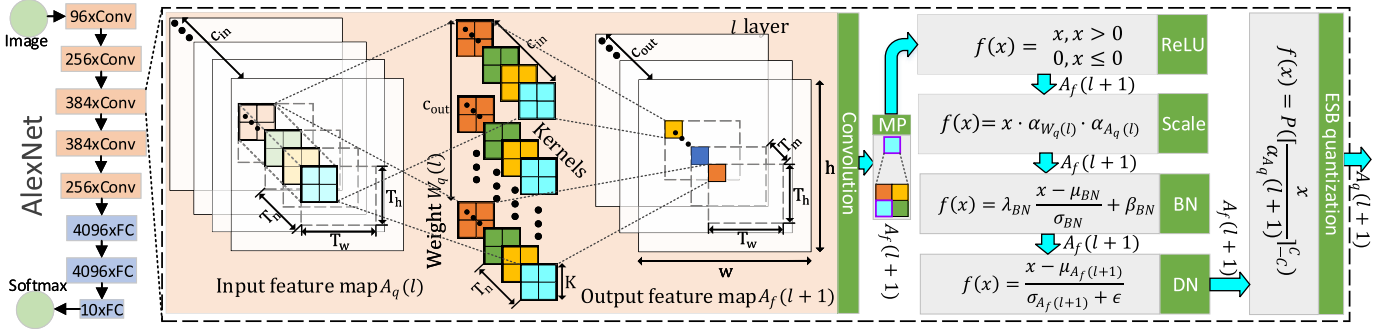


Fig. 9. Computing logic of a convolution layer in AlexNet.

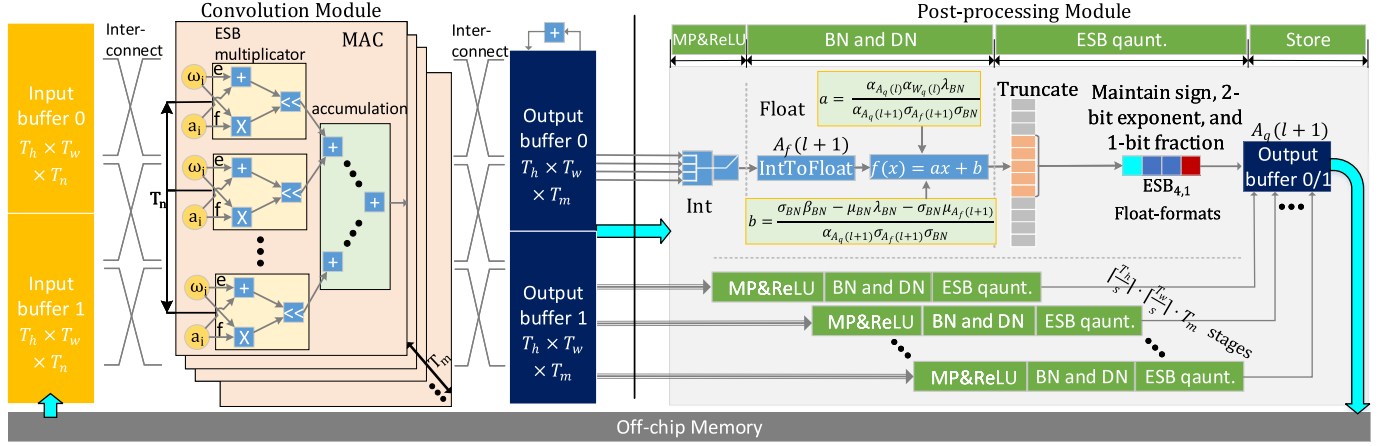


Fig. 10. Hardware architecture of convolution layer implementation.

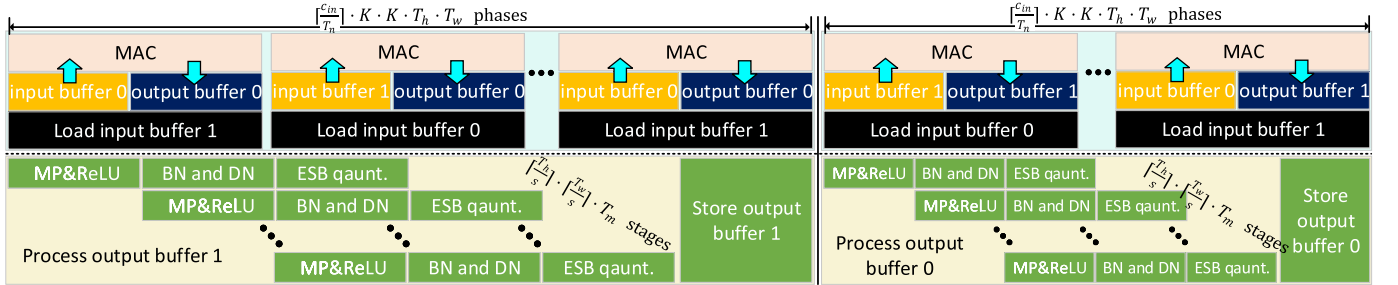


Fig. 11. Timing graph.

to reduce the computing overhead. The “ESB quantization” step, implemented with the truncation and projection as shown in (4), quantizes elements of $A_f(l+1)$ into the ESB float format and then stores them into $A_q(l+1)$. We design a pipeline to conduct the three steps. $A_q(l+1)$ is written to output buffer 0 or 1 and finally stored in the off-chip memory.

4.4 Timing Graph

Using the ESB pipeline design, we can compute the convolution module and post-processing module in parallel to fully utilize hardware resources and accelerate DNN inference. The timing graph of the accelerator is shown in Fig. 11. In the convolution module, we utilize dual ping-pong buffers for input prefetching to hide the data transmission time when executing MAC in the convolution module. The dual output buffers facilitate the execution of two

modules in parallel. When the convolution module writes data into the output buffer 0, the post-processing module deals with another output buffer 1, and vice versa.

5 EVALUATION

To validate the unique features of ESB and demonstrate its accuracy and efficiency benefits, we conducted extensive experiments on representative image classification tasks and implemented ESB quantization on FPGAs.

5.1 Experimental Setting

In the experiments, all the weights and activation of all layers in DNNs, including the first and last layers, are quantized by ESB with the same settings. We use Keras v2.2.4 and TensorFlow v1.13.1 as the DNN training frameworks. The layers in Keras are reconstructed as described in Section 3.5 to perform ESB quantization in the image classification task.

TABLE 1
Results of ESB Across Various Configurations

W/A	k	Name	Alias	DDA	α^*	Datasets & Models			
						Mnist LeNet5	Cifar10 VGG-like	ImageNet(Top1/Top5)	
								AlexNet	ResNet18
32/32	-	Referenced	-	0	-	99.40	93.49	60.01/81.90†	69.60/89.24†
2/2	0	ESB(2,0)	Ternary	0.1902	1.2240	99.40	90.55	58.47/80.58	67.29/87.23
3/3	0	ESB(3,0)	PoT	0.0476	0.5181	99.45	92.55	60.73/82.53	69.29/88.79
3/3	1	ESB(3,1)	Fixed	0.0469	1.3015	99.61	92.76	61.52/82.96	69.82/89.10
4/4	0	ESB(4,0)	PoT	0.0384	0.0381	99.61	92.54	61.02/82.38	69.82/89.04
4/4	1	ESB(4,1)	ESB	0.0127	0.4871	99.53	93.41	62.19/83.37	70.41/89.43
4/4	2	ESB(4,2)	Fixed	0.0129	1.4136	99.52	93.30	61.45/82.64	70.36/89.29
5/5	1	ESB(5,1)	ESB	0.0106	0.0391	99.55	93.44	61.68/83.08	70.41/89.57
5/5	2	ESB(5,2)	ESB	0.0033	0.4828	99.53	93.43	62.42/83.52	71.06/89.48
5/5	3	ESB(5,3)	Fixed	0.0037	1.5460	99.54	93.26	61.72/82.93	70.60/89.53
6/6	2	ESB(6,2)	ESB	0.0028	0.0406	99.52	93.53	62.07/83.31	70.30/89.58
6/6	3	ESB(6,3)	ESB	0.0008	0.4997	99.53	93.49	62.39/83.58	70.99/89.55
6/6	4	ESB(6,4)	Fixed	0.0011	1.6878	99.52	93.40	61.91/83.23	70.30/89.37
7/7	3	ESB(7,3)	ESB	0.0007	0.0409	99.52	93.29	61.99/83.32	71.06/89.55
7/7	4	ESB(7,4)	ESB	0.0002	0.5247	99.54	93.29	62.26/83.25	70.98/89.58
7/7	5	ESB(7,5)	Fixed	0.0003	1.8324	99.53	93.32	61.84/83.20	70.91/89.62
8/8	4	ESB(8,4)	ESB	0.0002	0.0412	99.52	93.47	62.22/83.40	71.03/89.78
8/8	5	ESB(8,5)	ESB	0.0001	0.5527	99.53	93.51	62.94/84.15	70.96/89.74
8/8	6	ESB(8,6)	Fixed	0.0001	1.9757	99.55	93.36	61.63/82.96	70.96/89.70

[†]W/A denotes the given bits for weight quantization and activation quantization, respectively. Name is abbreviated. Alias refers to special cases of ESB; Fixed denotes the linear quantization. DDA is the distribution difference estimation, and the bold results represent the minimal DDA under the corresponding given bits and significant bits. α^* indicates a desirable solution. The configurations of ESB(b, k) $b - k > 4$ are not recommended because they are similar to configurations of ESB($b - 1, k$) in terms of accuracy but consume more bits. † The results refer to [9].

Datasets and models. Five DNN models on three data sets are used in our experiments. We utilize LeNet5 with 32C5-BN-MP2-64C5-BN-MP2-512FC-10Softmax on MNIST and VGG-like with 64C3-BN-64C3-BN-MP2-128C3-BN-128C3-BN-MP2-256C3-BN-256C3-BN-MP2-1024FC-10-Softmax on Cifar10. Here, C is convolution, BN is batch normalization, MP indicates maximum pooling, and FC denotes full connection. Data augmentation in DSN [39] is applied for VGG-

like training. We also conduct experiments over two conventional networks, AlexNet [2] and ResNet18 [1], and a light-weight network MobileNetV2 [40] on ImageNet [41]. We use the BN [38] layer to replace the original local response normalization (LRN) layer in AlexNet for a fair comparison with ENN [18], LQ-Nets [30], and QIL [19]. We adopt the open data augmentations for ImageNet as in TensorFlow: resize images into a size of 224×224 following a

TABLE 2
Detailed Settings of Quantization Methods Collected From the Literature

Methods		Format	Weights			Activation			FConv	LFC
			Bits	SFB	SFN	Bits	SFB	SFN		
Linear (Fixed-point)	μ L2Q [8]	Fixed	2,4,8	32	1	32	-	-	Y	8
	VecQ [9]	Fixed	2,4,8	32	1	8	32	1	Y	Y
	TSQ [26]	Fixed	2	32	$c \uparrow$	2	32	2	N	N
	BCGD [20]	Fixed	4	32	1	4	32	1	Y	Y
	HAQ [28]	Fixed	f2,f3,f4§	32	1	32,f4	-/32	-/1	8	Y
	DSQ [7]	Fixed	2,3,4	32	1	2,3,4	32	1	N	N
	PACT [27]	Fixed	2,3,4,5	32	1	2,3,4,5	32	1	N	N
	Dorefa-Net [6]	Fixed	2,3,4,5	32	1	2,3,4,5	32	1	N	N
	QIL [19]	Fixed	2,3,4,5	32	1	2,3,4,5	32	1	N	N
Nonlinear	INQ [17]	PoT	2,3,4,5	32	1	32	-	-	-	-
	ENN [18]	PoT	2,3	32	1	32	-	-	-	-
	ABC-Net [21]	-	3,5	32	{3,5} $\ast c$	3,5	32	3,5	-	-
	LQ-Nets [30]	-	2,3,4	32	{2,3,4} $\ast c$	2,3,4	32	2,3,4	N	N
	AutoQ [31]	-	f2,f3,f4	32	{f2,f3,f4} $\ast c$	f2,f3,f4	32	f2,f3,f4	-	-
	APoT [15]	-	2,3,5	32	1	2,3,5	32	1	N	N
	ESB	ESB	2-8	32	1	2-8	32	1	Y	Y
	ESB*	ESB	2-5	32	1	2-5	32	1	N	N

[†]SFB is scaling-factor bitwidth, and SFN is the number of the scaling factors. FConv represents the first convolutional layer, LFC represents the last fully connected layer, and Y or N indicates whether quantization was done. Fixed refers to the fixed-point format. † c denotes the number of convolutional kernels. § fx denotes there are x average bits for the mixed-precision schemes.

random horizontal flip. Then, we subtract the mean of each channel for AlexNet and ResNet18, and scale pixels into the range of $[-1, 1]$ for MobileNetV2.

Model Training. We take the open pre-trained models from [9] to initialize the quantized models and then fine-tune them to converge. For experiments with fewer bits than 4, we train the quantized models for a total number of 60 epochs with an initial learning rate (LR) of 1.0. LR is decayed to 0.1, 0.01, and 0.001 at 50, 55, and 57 epochs, respectively. For experiments with given bits of more than 4, we train the quantized models for a total number of 10 epochs with an initial LR of 0.1, and decay the LR to 0.01 and 0.001 at 7 and 9 epochs, respectively.

Hardware Platform. Model training and evaluation are based on a *Super Micro 4028GR-TR* server equipped with two 12-core Intel CPUs (Xeon E5-2678 v3), an Nvidia RTX 2080ti GPU, and 64GB memory. The accelerator system of the ESB is implemented on the Xilinx ZCU102 FPGA platform, which is equipped with a dual-core ARM Cortex-A53 and based on Xilinx's 16nm FinFET+ programmable logic fabric with 274.08K LUTs, 548.16K FFs, 912 BRAM blocks with 32.1Mb, and 2520 DSPs. Vivado 2020.1 is used for system verification and ESB simulation.

5.2 ESB Accuracy Across Various Bits

We first focus on the investigation of the inference accuracy achieved by ESB across various bits. The notation with the value pair for ESB, such as $ESB(b, k)$, indicates that there are b given bits and k significant bits among b bits in the ESB scheme. As mentioned in Section 3.2, linear quantization ($k = b - 2$), PoT ($k = 0$), and Ternary ($k = 0$ and $b = 2$) are special cases of ESB. A series of experimental results are shown in Table 1 to verify the effect of elastic significant bits on the accuracy.

ESB can retain high accuracy across all datasets and DNN models. For example, when pair (b, k) has the values of (3,1), (6,2), (8,5), and (7,3), the accuracy results of ESB can reach 99.61%, 93.53%, 62.94%(Top1), and 71.06%(Top1), for LeNet5, VGG-like, AlexNet, and ResNet18, respectively, which outperform the referenced pre-trained models by 0.21%, 0.04%, 2.93%, and 1.46%, respectively. When the conditions $b - k = 3$ ($b > 3$) are satisfied, ESB has the minimal DDAs while maintaining relatively higher accuracy, such as $ESB(4,1)$ and $ESB(5,2)$. In particular, under three given bits, the results of ESB outperform referenced pre-trained AlexNet and ResNet18 by 1.51% and 0.22%, respectively, in terms of accuracy. Our proposed DDA calculates the distribution of quantized values of ESB to fit the bell-shaped distribution well, thus maintaining the features of the original values and reducing the accuracy degradation. In other words, the minimal DDA result corresponds to the best quantizing strategy of the ESB for different given bits.

5.3 Comparison With State-of-the-Art Methods

We then compare the ESB with state-of-the-art methods under the same bit configurations. We perform comparisons from two aspects: DNN inference efficiency and accuracy.

For efficiency, we obtain some observations through the quantization settings, as shown in Table 2. First, most of the methods, such as APoT [15], have to reserve full precision

TABLE 3
Comparison With the State-of-the-Art Approaches

Methods	W/A	AlexNet		ResNet18	
		Top1	Top5	Top1	Top5
INQ	2/32	-	-	66.02	87.13
ENN	2/32	58.20	80.60	67.00	87.50
μ L2Q	2/32	-	-	65.60	86.12
VecQ	2/8	58.04	80.21	67.91	88.30
AutoQ	f2/f3	-	-	67.40	88.18
TSQ	2/2	58.00	80.50	-	-
DSQ	2/2	-	-	65.17	-
PACT	2/2	55.00	77.70	64.40	85.60
Dorefa-Net	2/2	46.40	76.80	62.60	84.40
LQ-Nets	2/2	57.40	80.10	64.90	85.90
QIL	2/2	58.10	-	65.70	-
APoT	2/2	-	-	67.10	87.20
$ESB(2,0)$	2/2	58.47	80.58	67.29	87.23
$ESB^*(2,0)$	2/2	58.73	81.16	68.13	88.10
INQ	5/32	57.39	80.46	68.98	89.10
PACT	5/5	55.70	77.80	69.80	89.30
Dorefa-Net	5/5	45.10	77.90	68.40	88.30
QIL	5/5	61.90	-	70.40	-
ABC-Net	5/5	-	-	65.00	85.90
APoT	5/5	-	-	70.90	89.70
$ESB(5,2)$	5/5	62.42	83.52	71.06	89.48
$ESB^*(5,2)$	5/5	62.89	83.91	71.16	89.73
μ L2Q	8/32	-	-	65.52	86.36
VecQ	8/8	61.60	83.66	69.86	88.90
$ESB(8,5)$	8/8	62.94	84.15	70.96	89.74

Methods	W/A	AlexNet		ResNet18	
		Top1	Top5	Top1	Top5
Referenced	32/32	60.01	81.90	69.60	89.24
INQ	3/32	-	-	68.08	88.36
ENN2 §	3/32	59.20	81.80	67.50	87.90
ENN4 §	3/32	60.00	82.40	68.00	88.30
AutoQ	f3/f4	-	-	69.78	88.38
DSQ	3/3	-	-	68.66	-
PACT	3/3	55.60	78.00	68.10	88.20
Dorefa-Net	3/3	45.00	77.80	67.50	87.60
ABC-Net	3/3	-	-	61.00	83.20
LQ-Nets	3/3	-	-	68.20	87.90
QIL	3/3	61.30	-	69.20	-
APoT	3/3	-	-	69.70	88.90
$ESB(3,1)$	3/3	61.52	82.96	69.82	89.10
$ESB^*(3,1)$	3/3	61.99	83.46	70.14	89.36
INQ	4/32	-	-	68.89	89.01
μ L2Q	4/32	-	-	65.92	86.72
VecQ	4/8	61.22	83.24	68.41	88.76
DSQ	4/4	-	-	69.56	-
PACT	4/4	55.70	78.00	69.20	89.00
Dorefa-Net	4/4	45.10	77.50	68.10	88.10
LQ-Nets	4/4	-	-	69.30	88.80
QIL	4/4	62.00	-	70.10	-
BCGD	4/4	-	-	67.36	87.76
$ESB(4,1)$	4/4	62.19	83.37	70.41	89.43
$ESB^*(4,1)$	4/4	62.81	84.02	71.11	89.62

§ENN2 indicates one-bit shift results denoted as $\{-2, +2\}$, and ENN4 represents two-bit shift denoted as $\{-4, +4\}$ in ENN [18].

values for the first and last layers in DNNs; otherwise, they cannot prevent decreasing accuracy. Methods including μ L2Q [8], INQ [17], and ENN [18] are forced to use 32-bit floating-point activation to maintain accuracy. Second, some methods [21], [26], [30], [31] attempt to use

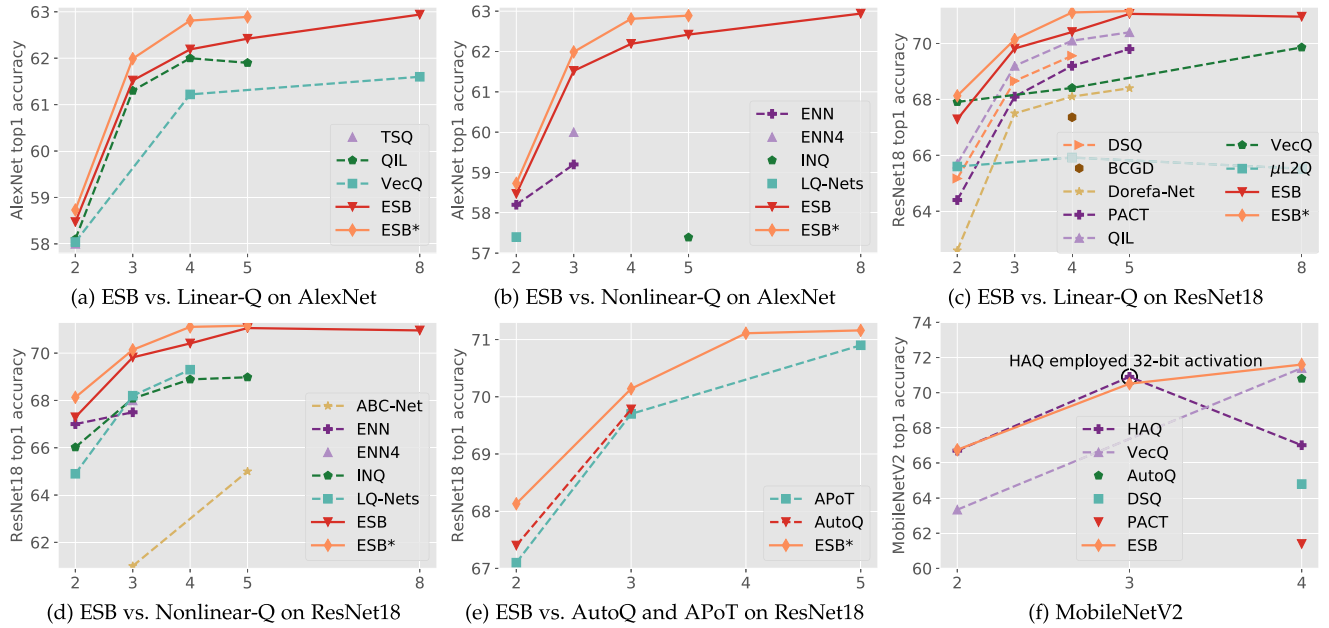


Fig. 12. Visual accuracy comparisons with state-of-the-art methods on different models. Linear-Q denotes linear quantization and Nonlinear-Q refers to nonlinear quantization. The x -axis represents bitwidth.

convolutional kernel quantization to replace weight quantization, and then introduce additional full precision scaling factors. The number of scaling factors of these methods is related to the number of convolutional kernels, as shown in Table 2, which can significantly increase the computing overhead.

In summary, ESB quantizes weights and activation across all layers with elastic significant bits and only introduces one shared full precision scaling factor for all elements of weights and activation. In this manner, ESB eliminates unnecessary computations and reduces the computational scale, thus improving efficiency. Moreover, in order to ensure a fair comparison with the methods that utilize the full precision weights/activation for the first and last layers, we perform additional experiments without quantizing the first and last layers. The results are denoted as *ESB**.

For accuracy comparison, we survey and replicate abundant existing methods as our comparison targets and list all the quantitative results in Table 3. Besides, to highlight the accuracy improvements of ESB, we provide the accuracy comparison by different quantization categories (Linear-Q and Nonlinear-Q) in Fig. 12. ESB consistently outperforms the state-of-the-art approaches across all given bits and improves the top 1 results of AlexNet and ResNet18 by 4.78% and 1.92%, respectively, on average. Even by quantizing the weights and activation for the first and last layers, ESB still outperforms APoT by 0.16% for accuracy on average. APoT performs a simple sum of several PoTs, which leads to the deviation of the target bell-shaped distribution far away. ESB also has higher accuracy compared with the approaches using full precision activation such as INQ and ENN, with improvements of 1.65% and 1.48% over ResNet18, respectively. In addition, ESB is suitable for 2-bit quantization and is better than 2-bit schemes such as TSQ, DSQ, and APoT in terms of accuracy. Furthermore, ESB* consistently achieves the best accuracy and outperforms all state-of-the-art works over AlexNet and

ResNet18, respectively. Even under two given bits, compared with APoT, ESB* can improve accuracy by up to 1.03%. On average, ESB* outperforms APoT by 0.58% in terms of accuracy.

The reason ESB can achieve high accuracy is that ESB fits the distribution well with the ESB and DDA. Significant bits provide a fine-grain control of the density of quantized values for various distributions of DNN weights/activation. Furthermore, DDA can minimize the difference between the ESB quantized value distribution and the original distribution to facilitate ESB to seek the best distribution. Consequently, fine control granularity inspires ESB to achieve accuracy gains.

5.4 ESB on Light-Weight Model

We next evaluate the ESB on the popular light-weight model, MobileNetV2 [40], to verify the superiority of ESB. The important results of the ESB are presented in Table 4 and Fig. 12f. We would like to emphasize that the ESB quantizes the weights and activation for all layers of MobileNetV2. ESB outperforms DSQ [7] and VecQ [9], and achieves an average accuracy improvement of 3.56% compared with state-of-the-art methods. Lowering to 2 bits for quantization, ESB can maintain an accuracy of 66.75%, which is similar to the accuracy of HAQ with mixed-precision weights and full precision activation. Experiments on MobileNetV2 demonstrate that ESB quantization is an effective strategy to fit the distributions of weights or activation by controlling the number of significant bits of quantized values, so as to mitigate the accuracy degradation of low-bitwidth quantization.

5.5 Hardware Evaluation

In this subsection, we implement ESB accelerator on the FPGA platform to verify the hardware benefits of the ESB. The objectives of the hardware evaluation are twofold: 1) exploring the hardware resource utilization of ESB with

TABLE 4
Results on Light-Weight MobileNetV2 [40]

Methods	W/A	MobileNetV2	
		Top1	Top5
Referenced	32/32	71.30	90.10
HAQ	f2/32	66.75	87.32
VecQ	2/8	63.34	84.42
ESB(2,0)	2/2	66.75	87.14
HAQ	f3/32	70.90	89.76
ESB(3,1)	3/3	70.51	89.46

Methods	W/A	MobileNetV2	
		Top1	Top5
VecQ	4/8	71.40	90.41
HAQ	f4/f4	67.01	87.46
AutoQ	f4/f4	70.80	90.33
DSQ	4/4	64.80	-
PACT	4/4	61.39	83.72
ESB(4,1)	4/4	71.60	90.23

[†]The results for Referenced refer to [9], and the results of PACT are cited from [28].

different configurations; 2) providing insight into energy efficiency and presenting comparisons among ESB accelerator, state-of-the-art designs, and CPU/GPU platforms.

Resource Utilization. In general, there are limited computing resources such as LUTs, FFs, BRAMs, and DSPs in FPGAs. Hardware implementations usually need to fully utilize the resources to achieve a high concurrency and clock frequency. In our experiments, we freeze the clock frequency to 145 MHz. For a fair comparison, we employ the LUT instead of the DSP to implement fraction multiplication. The resource utilization and concurrency under various ESB implementations are listed in Table 5.

First, we show the effect of how significant bits affect the MAC, including the multiplier and accumulator, at the bit-wise level. We observe that fewer significant bits of fraction, that is, a smaller value of k , can reduce the resource consumption of the multiplier. With the same number of given bits, the consumption of LUTs for one multiplier can

decrease by 6.3 on an average if k is decreased by 1. This is because fewer significant bits can reduce the consumption for fraction-multiplication, that is, $(\zeta_i + 0.f_i) * (\zeta_j + 0.f_j)$ in (15), thus consuming fewer computing resources. Nevertheless, the results also show that decreasing k can increase the resource consumption of the accumulator. Using the same number of given bits, the consumption of LUTs of the accumulator increases by 3.6 on average, if k is decreased by 1. This is because the relatively large exponents, that is, $e = e_i\zeta_i + e_j\zeta_j$ in (15), require a high output bitwidth and thus increase the resource budgets of the accumulator. In terms of MAC consumption, ESB can strike a better balance between the multiplier and accumulator, utilizing fewer resources than fixed-point (Fixed) and PoT.

Second, we verify that the elastic significant bits can improve the performance of the accelerator design. We take the achieved performance under per 1k LUTs for a quantitative evaluation. As shown in Fig. 13, the accelerator of ESB(2,0) exhibits the best performance with 10.59 GOPS under 1k LUTs, because it costs only 14 LUTs for MAC. For the other cases, the best results are 8.63, 5.27, 3.59, 2.70, 2.27, and 1.99 GOPS/kLUTs for configurations (3,1), (4,1), (5,1), (6,3), (7,5), and (8-5) across various quantizing bits, and the improvements are up to 11.19%, 18.68%, 10.77%, 12.96%, 42.59%, and 25.73% compared with the results of (3,0), (4,0), (5,3), (6,4), (7,3), and (8,4), respectively. By employing desirable number of significant bits in accelerator design, the performance within per 1k LUTs can improve by 20.32% on average without requiring more given bits.

Energy Efficiency. Energy efficiency is a critical criterion for evaluating a computing system, which is defined as the number of operations that can be completed within a unit of energy (J) [42]. Here, we define the energy efficiency as the ratio of peak performance (GOPS) to power (W). Fig. 14 presents the energy efficiency comparisons of our accelerator with the recent FPGA designs, including ELB-NN[43], Winograd [44], Caffeine [45], and Roofline-model-based method (RMB) [46], and CPU (Xeon E5-2678 v3)/GPU (RTX 2080 Ti) platforms.

TABLE 5
ESB Accelerator Resource Utilization and Performance Based on AlexNet

ESB(b, k)	Alias	MAC (#LUTs)			Resource Utilization (%)				Concurrency		Clock (MHz)	Power (W)	GOPS
		Mul.†	Acc.†	Mul.+Acc.	LUTs	FF	BRAM	DSP	Tn	Tm			
ESB(2,0)	Ternary	2	12	14	63.94	16.27	47.64	4.25	32	200	145	3.57	1856.00
ESB(3,0)	PoT	6	16	22	69.81	18.21	41.06	4.48	32	160	145	4.30	1484.80
ESB(3,1)	Fixed†	5	15	20	62.78	17.89	41.06	4.48	32	160	145	4.67	1484.80
ESB(4,0)	PoT	11	24	35	73.24	18.04	46.33	4.25	32	96	145	3.67	890.88
ESB(4,1)	ESB	11	19	30	61.71	16.33	30.54	4.25	32	96	145	4.08	890.88
ESB(4,2)	Fixed	19	17	36	70.86	15.21	30.54	4.25	32	96	145	4.02	890.88
ESB(5,1)	ESB	20	27	47	60.37	15.99	49.84	4.33	32	64	145	3.93	593.92
ESB(5,2)	ESB	26	21	47	60.95	14.44	49.84	4.33	32	64	145	3.89	593.92
ESB(5,3)	Fixed	41	19	60	66.87	14.17	39.31	4.33	32	64	145	3.79	593.92
ESB(6,2)	ESB	36	29	65	61.07	14.68	45.01	4.68	16	92	145	3.70	426.88
ESB(6,3)	ESB	45	23	68	57.64	13.89	45.01	4.68	16	92	145	4.78	426.88
ESB(6,4)	Fixed	55	21	76	65.11	13.67	45.01	4.68	16	92	145	4.65	426.88
ESB(7,3)	ESB	59	31	90	85.24	14.99	41.06	4.44	16	80	145	4.02	371.20
ESB(7,4)	ESB	57	25	82	60.88	13.86	41.06	4.44	16	80	145	3.73	371.20
ESB(7,5)	Fixed	66	23	89	59.78	14.19	41.06	4.44	16	80	145	4.12	371.20
ESB(8,4)	ESB	71	33	104	74.90	15.49	37.77	4.44	16	70	145	4.89	324.80
ESB(8,5)	ESB	69	27	96	59.57	14.55	37.77	4.44	16	70	145	3.99	324.80
ESB(8,6)	Fixed	86	25	111	63.17	14.35	37.77	4.44	16	70	145	3.61	324.80

†Mul. denotes the multiplier; Acc. represents the accumulator; Fixed refers to the fixed-point format of linear quantization.

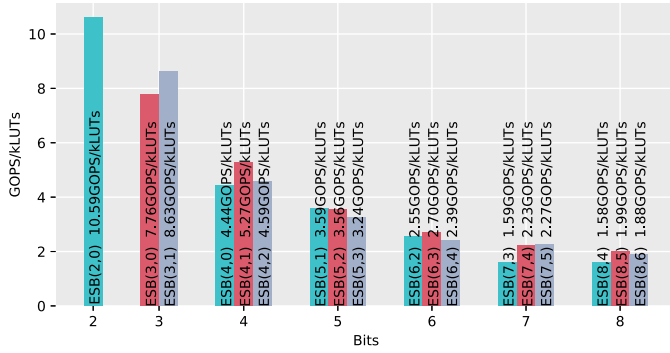


Fig. 13. Peak performance (GOPS) per 1k LUTs.

Compared with the ELB-NN [43] for the formats of 8-8218 and 4-8218 (extremely low bitwidth quantization and the least bit width is 1), ESB(2,0) achieves higher energy efficiency, with improvements of 71.90% and 9.73%, respectively. Compared with the accelerators using high-bitwidth formats, such as Fixed-8 [43], [47] and Fixed-16 [44], [45], or even float-point data type (FP32) [46], ESB(8,6) can achieve a higher energy efficiency, and the improvement is by $8\times$ on average, and up to $26\times$ (ESB(8,6) versus [46]). ESB accelerator also achieves a higher energy efficiency than CPU/GPU platforms across all bitwidths and configurations. Their energy efficiency increased by $22\times$ and $4\times$ on average compared with CPU and GPU, and up to $65\times$ and $11\times$ improvement at 2-bit, respectively. In addition, ESB accelerator can be implemented with various bitwidths and configurations, which can help us achieve a fine-grained trade-off between inference accuracy and computing efficiency. For example, ESB(3,0) is a reasonable choice under a tolerated accuracy of 60% (AlexNet Top1), which can achieve an energy efficiency of 345.30GOPS/W with an accuracy of 60.73%. ESB(3,1) with 317.94GOPS/W is a reasonable choice under a tolerated accuracy of 61%, while ESB(4,1) with 218.35GOPS/W is the desirable choice under a tolerated accuracy of 62%.

6 CONCLUSION

In this paper, we propose an ESB for DNN quantization to strike a good balance between accuracy and efficiency. We explore a new perspective to improve quantization accuracy and concurrently decrease computing resources by

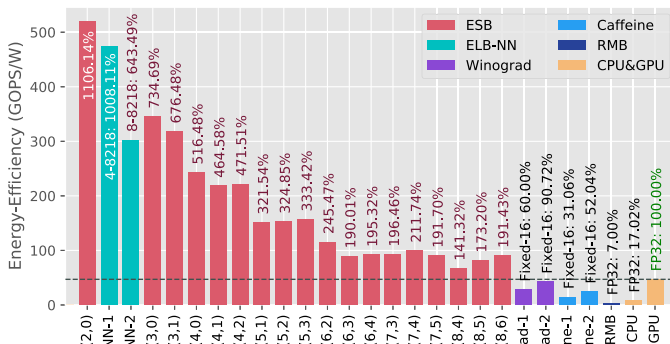


Fig. 14. Energy efficiency comparison among state-of-the-art designs and CPU/GPU.

controlling significant bits of quantized values. We propose the mathematical expression of ESB to unify various distributions of quantized values. The unified expression outlines the linear quantization and PoT quantization, and guides DNN compression from both accuracy and efficiency perspectives by fitting distribution using ESB. Moreover, we design an ESB projection operator involving shifting and rounding operations, which are cheap bit operations at the hardware level without complex conversions. Fewer significant bits representing weight/activation and cheap projection of quantization are what hardware desires because they meet hardware computing natures. In addition, the proposed DDA always finds a desirable distribution for ESB quantization to reduce accuracy degradation.

Experimental results show that ESB quantization outperforms state-of-the-art methods and achieves higher accuracy across all quantizing bitwidths. The hardware implementation also demonstrates that quantizing DNNs with elastic significant bits can achieve higher energy efficiency without increasing the quantizing bitwidth. In the future, we will further study the real performance benefits of our ESB on edge devices.

ACKNOWLEDGMENTS

Cheng Gong and Ye Lu contributed equally.

REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1106–1114.
- [3] W. Niu *et al.*, "PatDNN: Achieving real-time DNN execution on mobile devices with pattern-based weight pruning," in *Proc. Architect. Support Program. Languages Operating Syst.*, 2020, p. 907–922.
- [4] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: Imagenet classification using binary convolutional neural networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 525–542.
- [5] Y. Chen *et al.*, "T-DLA: An open-source deep learning accelerator for ternarized dnn models on embedded FPGA," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, 2019, pp. 13–18.
- [6] S. Zhou, Z. Ni, X. Zhou, H. Wen, Y. Wu, and Y. Zou, "DoReFa-Net: Training low bitwidth convolutional neural networks with low bitwidth gradients," *CoRR*, vol. abs/1606.06160, 2016.
- [7] R. Gong *et al.*, "Differentiable soft quantization: Bridging full-precision and low-bit neural networks," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 4851–4860.
- [8] C. Gong *et al.*, "μ12Q: An ultra-low loss quantization method for DNN compression," in *Proc. Int. Joint Conf. Neural Netw.*, 2019, pp. 1–8.
- [9] C. Gong, Y. Chen, Y. Lu, T. Li, C. Hao, and D. Chen, "VecQ: Minimal loss DNN model compression with vectorized weight quantization," *IEEE Trans. Comput.*, vol. 70, no. 5, pp. 696–710, May 2021.
- [10] R. Banner, Y. Nahshan, and D. Soudry, "Post training 4-bit quantization of convolutional networks for rapid-deployment," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 7950–7958.
- [11] J. Faraone, N. J. Fraser, M. Blott, and P. H. W. Leong, "SYQ: Learning symmetric quantization for efficient deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4300–4309.
- [12] B. Jacob *et al.*, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 2704–2713.
- [13] D. Miyashita, E. H. Lee, and B. Murmann, "Convolutional neural networks using logarithmic data representation," *CoRR*, vol. abs/1603.01025, 2016.

- [14] J. Faraone *et al.*, "AddNet: Deep neural networks using FPGA-optimized multipliers," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 28, no. 1, pp. 115–128, Jan. 2020.
- [15] Y. Li, X. Dong, and W. Wang, "Additive powers-of-two quantization: An efficient non-uniform discretization for neural networks," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [16] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding," in *Proc. Int. Conf. Learn. Representations*, 2016.
- [17] A. Zhou, A. Yao, Y. Guo, L. Xu, and Y. Chen, "Incremental network quantization: Towards lossless CNNs with low-precision weights," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [18] C. Leng, Z. Dou, H. Li, S. Zhu, and R. Jin, "Extremely low bit neural network: Squeeze the last bit out with ADMM," in *AAAI*, S. A. McIlraith and K. Q. Weinberger, Eds. Palo Alto, CA, USA: AAAI Press, 2018, pp. 3466–3473.
- [19] S. Jung *et al.*, "Learning to quantize deep networks by optimizing quantization intervals with task loss," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4350–4359.
- [20] P. Yin, S. Zhang, J. Lyu, S. J. Osher, Y. Qi, and J. Xin, "Blended coarse gradient descent for full quantization of deep neural networks," *CoRR*, vol. abs/1808.05240, 2018.
- [21] X. Lin, C. Zhao, and W. Pan, "Towards accurate binary convolutional neural network," in *Proc. Conf. Neural Inf. Process. Syst.*, 2017, pp. 345–353.
- [22] X. Fan *et al.*, "Effect of image noise on the classification of skin lesions using deep convolutional neural networks," *Tsinghua Sci. Technol.*, vol. 25, no. 3, pp. 425–434, 2020.
- [23] Q. Cao, W. Zhang, and Y. Zhu, "Deep learning-based classification of the polar emotions of "moe"-style cartoon pictures," *Tsinghua Sci. Technol.*, vol. 26, no. 3, pp. 275–286, 2021.
- [24] W. Zhong, N. Yu, and C. Ai, "Applying Big Data based deep learning system to intrusion detection," *Big Data Min. Anal.*, vol. 3, no. 3, pp. 181–195, 2020.
- [25] Z. Kai and Z. Tao, "Deep reinforcement learning based mobile robot navigation: A review," *Tsinghua Sci. Technol.*, vol. 26, no. 5, pp. 674–691, 2021.
- [26] P. Wang, Q. Hu, Y. Zhang, C. Zhang, Y. Liu, and J. Cheng, "Two-step quantization for low-bit neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4376–4384.
- [27] J. Choi, Z. Wang, S. Venkataramani, P. I. Chuang, V. Srinivasan, and K. Gopalakrishnan, "PACT: Parameterized clipping activation for quantized neural networks," *CoRR*, vol. abs/1805.06085, 2018.
- [28] K. Wang, Z. Liu, Y. Lin, J. Lin, and S. Han, "HAQ: Hardware-aware automated quantization with mixed precision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8612–8620.
- [29] M. Ghasemzadeh, M. Samragh, and F. Koushanfar, "ReBNet: Residual binarized neural network," in *Proc. 26th IEEE Annu. Int. Symp. Field-Programmable Custom Comput. Mach.*, 2018, pp. 57–64.
- [30] D. Zhang, J. Yang, D. Ye, and G. Hua, "LQ-Nets: Learned quantization for highly accurate and compact deep neural networks," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 373–390.
- [31] Q. Lou, F. Guo, M. Kim, L. Liu, and L. Jiang, "AutoQ: Automated kernel-wise neural network quantization," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [32] M. Drumond, T. Lin, M. Jaggi, and B. Falsafi, "End-to-end DNN training with block floating point arithmetic," *CoRR*, vol. abs/1804.01526, 2018.
- [33] U. Köster *et al.*, "Flexpoint: An adaptive numerical format for efficient training of deep neural networks," in *Proc. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1742–1752.
- [34] J. Gustafson and I. Yonemoto, "Beating floating point at its own game: Posit arithmetic," *Supercomputing Front. Innovations*, vol. 4, no. 2, 2017.
- [35] Z. Carmichael, H. F. Langroudi, C. Khazanov, J. Lillie, J. L. Gustafson, and D. Kudithipudi, "Deep positron: A deep neural network using the posit number system," in *Proc. Des. Autom. Test Europe Conf. Exhibit.*, 2019, pp. 1421–1426.
- [36] T. Tambe *et al.*, "Adaptivfloat: A floating-point based data type for resilient deep learning inference," *CoRR*, vol. abs/1909.13271, 2019.
- [37] W. Chen, J. Wilson, S. Tyree, K. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2285–2294.
- [38] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [39] C. Lee, S. Xie, P. W. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in *Proc. 18th Int. Conf. Artif. Intell. Statist.*, 2015, pp. 562–570.
- [40] M. Sandler, A. G. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4510–4520.
- [41] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.
- [42] K. Guo, S. Zeng, J. Yu, Y. Wang, and H. Yang, "[DL] A survey of FPGA-based neural network inference accelerators," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 12, no. 1, pp. 1–26, 2019.
- [43] J. Wang, Q. Lou, X. Zhang, C. Zhu, Y. Lin, and D. Chen, "Design flow of accelerating hybrid extremely low bit-width neural network in embedded FPGA," in *Proc. 28th Int. Conf. Field Programmable Logic Appl.*, 2018, pp. 163–1636.
- [44] L. Lu, Y. Liang, Q. Xiao, and S. Yan, "Evaluating fast algorithms for convolutional neural networks on FPGAs," in *Proc. IEEE 25th Annu. Int. Symp. Field-Programmable Custom Comput. Mach.*, 2017, pp. 101–108.
- [45] C. Zhang, G. Sun, Z. Fang, P. Zhou, P. Pan, and J. Cong, "Caffeine: Toward uniformed representation and acceleration for deep convolutional neural networks," *IEEE Trans. Comput-Aided Design Integr. Circuits Syst.*, vol. 38, no. 11, pp. 2072–2085, Nov. 2019.
- [46] C. Zhang, P. Li, G. Sun, Y. Guan, B. Xiao, and J. Cong, "Optimizing FPGA-based accelerator design for deep convolutional neural networks," in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays*, 2015, pp. 161–170.
- [47] M. Véstias, R. P. Duarte, J. T. de Sousa, and H. Neto, "Lite-CNN: A high-performance architecture to execute CNNs in low density FPGAs," in *Proc. 28th Int. Conf. Field Programmable Logic Appl.*, 2018, pp. 399–3993.



Cheng Gong received the BEng degree in computer science from Nankai University, in 2016. He is currently working toward the PhD degree at the College of Computer Science, Nankai University. His research interests include heterogeneous computing, machine learning and Internet of Things.



Ye Lu received the BS and PhD degrees from Nankai University, Tianjin, China, in 2010 and 2015, respectively. He is currently an associate professor with the College of Cyber Science, Nankai University now. His research interests include DNN FPGA accelerator, blockchain virtual machine, embedded system, Internet of Things.



Kunpeng Xie received the BEng degree in computer science from Nankai University, in 2019. He is currently working toward the PhD with the College of Computer, Intelligent computing system Lab, Nankai University. His research focuses on how computer CPU-GPU heterogeneous system accelerates the training of neural network models and to design effective CNN inference accelerator on FPGA.



Zongming Jin received the BS degree from Xidian University, Xian, China, in 2019. He is currently working toward the MS degree with Intelligent Computing System Lab, College of Cyber Science, Nankai University, Tianjin, China. His research interests include heterogeneous computing, embedded system and Internet of Things.



Tao Li received the PhD degree in computer science from Nankai University, China in 2007. He is currently a professor with the College of Computer Science, Nankai University. He is the member of IEEE Computer Society and the ACM, and the distinguished member of the CCF. His research interests include heterogeneous computing, machine learning and Internet of Things.



Yanzhi Wang (Member, IEEE) received the BS degree in electronic engineering from Tsinghua University, in 2009, and the PhD degree in computer engineering from the University of Southern California (USC), in 2014, under the supervision of Prof. Massoud Pedram. He received the Ming Hsieh Scholar Award (the highest honor in the EE Department of USC) for his PhD study. He is currently an assistant professor with the Department of Electrical and Computer Engineering, and Khoury College of Computer Science (Affiliated), Northeastern University. His research interests include model compression and platform-specific acceleration of deep learning applications.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**