

Hardware-Aware DNN Compression via Diverse Pruning and Mixed-Precision Quantization

Konstantinos Balaskas, Andreas Karatzas, Christos Sad, Kostas Siozios, *Senior Member, IEEE*, Iraklis Anagnostopoulos, *Member, IEEE*, Georgios Zervakis, Jörg Henkel, *Fellow, IEEE*

Abstract—Deep Neural Networks (DNNs) have shown significant advantages in a wide variety of domains. However, DNNs are becoming computationally intensive and energy hungry at an exponential pace, while at the same time, there is a vast demand for running sophisticated DNN-based services on resource constrained embedded devices. In this paper, we target energy-efficient inference on embedded DNN accelerators. To that end, we propose an automated framework to compress DNNs in a hardware-aware manner by jointly employing pruning and quantization. We explore, for the first time, per-layer fine- and coarse-grained pruning, in the same DNN architecture, in addition to low bit-width mixed-precision quantization for weights and activations. Reinforcement Learning (RL) is used to explore the associated design space and identify the pruning-quantization configuration so that the energy consumption is minimized whilst the prediction accuracy loss is retained at acceptable levels. Using our novel composite RL agent we are able to extract energy-efficient solutions without requiring retraining and/or fine tuning. Our extensive experimental evaluation over widely used DNNs and the CIFAR-10/100 and ImageNet datasets demonstrates that our framework achieves 39% average energy reduction for 1.7% average accuracy loss and outperforms significantly the state-of-the-art approaches.

Index Terms—Deep Neural Networks, DNN accelerators, DNN compression, Energy efficiency, Pruning, Quantization, Reinforcement Learning

1 INTRODUCTION

During the last decade, Deep Neural Networks (DNNs) have been established as the driving force in a wide range of application domains, such as object detection, speech recognition, virtual/augmented reality and more [1]. Modern DNN architectures are becoming increasingly more complex and deeper, with an ever increasing number of trainable parameters, striving to maximize the accuracy of modern applications. In order to keep up with the increased computational complexity, hardware DNN accelerators [2] form an integral part of computing systems. Such accelerators integrate thousands of multiply-accumulate (MAC) units, responsible for executing the required arithmetic operations [3]. Even though this large number of MAC units, operating in parallel, results in improved performance, it comes with the burden of vast energy consumption, thus

greatly affecting their integration in energy-constrained embedded systems [3].

In order to balance the trade-off between accuracy and energy consumption, several DNN compression techniques have been exploited. Noticeably, pruning and quantization have emerged as the most popular, mainly due to their effectiveness and simplicity. Pruning consists of intelligently sparsifying a dense DNN, by either removing connections (i.e., fine-grained pruning [4]–[6]) or regular structures (i.e., coarse-grained pruning [7]–[9]). Quantization [10], [11] on the other hand, lowers the arithmetic precision of operands (i.e., weights and activations), thus reducing the energy consumption of the DNN accelerator as well. However, in both cases the DNN accuracy is affected and thus, methods to mitigate the associated loss are required.

Earlier works on DNN pruning attempted to maximize network sparsity or reduce MAC operations, using either a fine- or coarse-grained approach. Though, exploiting pruning at both granularities at the same time is an uncharted territory for DNN compression. Moreover, indirect optimization metrics (e.g., MAC count) proved suboptimal in maximizing or even achieving latency/power gains on hardware [12], mandating the need for hardware-aware compression frameworks with direct platform feedback [13]. Still, combining the benefits of pruning and quantization remains mostly unexplored, mainly due to the highly increased complexity of their joint application [14]. To address this limitation, related works attempted to combine pruning with quantization by following a learning-based approach (e.g., with reinforcement learning (RL)) [15] and showed promising results. However, in such works, fine-tuning and/or long-term retraining plays an essential role in improving upon the DNN accuracy. Though, fine-tuning

- K. Balaskas, C. Sad and K. Siozios are with the Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece. K. Balaskas is also with the Chair for Embedded Systems at Karlsruhe Institute of Technology, Karlsruhe 76131, Germany.
- A. Karatzas and I. Anagnostopoulos are with the School of Electrical, Computer and Biomedical Engineering, Southern Illinois University Carbondale, Carbondale, IL 62901 USA.
- G. Zervakis is with the Dept. of Computer Engineering & Informatics, University of Patras, Patras 26504, Greece. This research was done when he was with the Karlsruhe Institute of Technology.
- J. Henkel are with the Chair for Embedded Systems, Karlsruhe Institute of Technology, Karlsruhe 76131, Germany.

Corresponding author: Konstantinos Balaskas (balaskas@kit.edu).

This work is supported in parts by grant NSF CCF 2324854, by the German Research Foundation (DFG) project “ACCROSS” HE 2343/16-1 under the grant 428566201 and by the E.C. Funded Program “SERRANO” under H2020 Grant 101017168.

Manuscript received October 13, 2022, revised July 31, 2023 and November 15, 2023.

during the optimization phase and/or retraining afterwards, tremendously increases the required execution time, especially when dealing with large classification datasets (e.g., ImageNet). In addition, such approaches may be infeasible due to proprietary or private training data [16], and/or when targeting resource constrained embedded devices that do not support training. For example, considering the latter, assume a resource constrained embedded system that comprises an inference DNN accelerator. Deploying a hardware-aware compression optimization that requires fine-tuning/retraining would be infeasible in that scenario. Hence, although on-device optimization would enable hardware-aware compression and usage of private local data, existing state-of-the-art techniques [15], [17], [18], due to their inherent requirement for fine-tuning, cannot be deployed on resource constrained embedded systems.

In this work, we address the aforementioned limitations and propose a learning-based compression framework for energy-efficient DNN inference on accelerator platforms. To achieve this, we apply layer-wise pruning and mixed-precision quantization. Our framework explores the full spectrum of the joint pruning and quantization design space in a hardware-aware manner, without any retraining/fine-tuning (i.e., one-shot). Our approach is generic to the pruning criterion, meaning that an optimal (different) pruning algorithm is selected per layer, from a predefined set of diverse fine- and coarse-grained pruning techniques. To handle the huge design space, we utilize RL to automatically prune and quantize an input DNN. We propose a novel composite RL agent, able to learn a compression policy which optimally balances the accuracy/energy trade-off. We evaluate our framework over state-of-the-art image classification models, on the CIFAR-10, CIFAR-100, and ImageNet datasets. Our experimental results demonstrate that the compressed DNNs generated by our framework feature small accuracy loss and up to 57% energy gains.

Our main contributions in this work are the following:

- (1) We propose an automated framework for energy-efficient DNN inference, which applies for the first-time one-shot compression in a learning-based manner and with hardware-aware techniques.
- (2) We boost the pruning efficiency by considering a diverse set of pruning algorithms and combining, for the first time, fine- with coarse-grained pruning.
- (3) We propose a novel composite RL agent, able to learn the optimal pruning and quantization profile for each layer.

The rest of the paper is organized as follows: Section 2 provides an overview of the existing literature on DNN compression-related topics. Section 3 serves as the motivation for our work, by showcasing via examples the benefits of including diverse pruning algorithms (instead of monolithic ones) and mixed-precision quantization (instead of uniform). In Section 4, the core idea of our work is presented, detailing our compression technique, the implemented composite RL agent and finally, our energy estimation model. Section 5 presents the evaluation setup and results, focusing on comparisons against the state of the art and the efficacy of our exploration mechanism. Section 6 concludes this paper.

TABLE 1: Taxonomy of the differentiating characteristics between existing state-of-the-art works and our framework.

	AMC [15]	HAQ [17]	ASQJ [24]	Ours
Fine-grained Pruning	-	-	✓	✓
Coarse-grained Pruning	✓	-	-	✓
Mixed-Precision Quantization	-	✓	✓	✓
Hardware feedback	✓	✓	-	✓
Without retraining	-	-	-	✓

2 RELATED WORK

Several techniques for DNN pruning and quantization have been proposed throughout the years, aiming to reduce the memory footprint while retaining acceptable prediction accuracy of DNNs. Earlier works on pruning [4], [5] focused on fine-grained removal of connections (i.e., weight pruning) based on sensitivity analyses. In [19], authors contemplated that randomly-initialized dense DNNs contain subnetworks, able to match the accuracy of their originator after retrained in isolation. [12] presents a hardware-aware pruning framework to adapt a DNN to a given mobile platform, driven by direct hardware metrics to meet constrained budget requirements. In [13], fine-grained pruning is guided by energy estimations from arbitrary hardware platforms for energy-efficient DNN inference. In [20], a methodology to prune fine-grained patterns inside coarse grain structures is proposed, which, combined with compiler-level optimization techniques, leads to inference speedup. However, quantization is not considered by the above.

Focusing on post-training quantization, a three part methodology was proposed in [21], composed of activation clipping, per-channel bit allocation, and weight bias correction, targeting 4-bits precision. However, not all layers are quantized. Authors in [22] duplicate channels containing outliers and halve their outputs, decreasing, thus, the quantization error and preserving network accuracy. In [23], the quantization step and scaling factors for DNN weights are discovered through minimization of a composite vector loss instead of the quantization error. However, the impact on hardware efficiency is not considered.

Improving on rule-based methods, RL has been employed to discover optimal compression solutions in a learning-based manner [15], [17], [25]. Automatic model compression [15] aims to provide layer-wise sparsity targets by utilizing a Deep Deterministic Policy Gradient (DDPG) agent in several hardware-emulating constraint scenarios. In [17], on the other hand, an RL-based framework is proposed to find layer-wise optimal mixed-precisions, based on hardware-aware feedback from arbitrary platforms. Both approaches consider only a single compression technique (pruning or quantization) to limit the associated design space explored by the RL agent.

Pruning and quantization are orthogonal to each other and can be applied synergistically for larger energy and memory gains. Han et. al [26] were the first to combine weight pruning with quantization (by means of clustering and weight sharing) and Huffman coding to further compress DNNs. In [24], an automated framework for jointly pruning and quantization was based on alternating direc-




Fig. 1: Evaluation of accuracy loss (left y-axis) and energy gain (right y-axis) of different sparsity rates (x-axis) for (a) VGG16, (b) ResNet50 and (c) MobileNetV2 on CIFAR-10. Color separates fine-grain (blue) and coarse-grain pruning (green). Shape separates between accuracy loss (triangles) and energy gain (diamonds).

tion method of multipliers (ADMM). In [14], subnetworks from an once-for-all network [27] (found via architectural search) were sampled, quantized with mixed-precision and fed to a heuristic optimization to determine optimal compression rate and precision. A differentiable approach is proposed in [28], as a joint loss function is constructed to achieve a balance between compression techniques whilst preserving the accuracy during training. In [29], pruning and quantization are applied via clipping and quantizing the remaining weights during training, in a single stage. Similarly, [30] proposes a technique to scale each gradient based on the weight vector position, acting as a regularizer for pruning/quantization-aware training. A weight representation scheme is proposed in [31], tailored for irregular sparse matrices in a hardware friendly manner (i.e., exploiting XOR gates). PQK [32] is a distillation-based pruning framework which utilizes (previously unimportant) pruned and quantized weights in the distillation phase, instead of pretrained teacher models.

Some related works apply pruning and quantization in one-shot, meaning that compression is applied without access to training data. In [18], authors use an analytical Lagrangian-based approximation model to find layer-wise pruning masks and quantization steps, followed by fine-tuning. In [33], an exact mathematical model is used for post-training one-shot compression, followed by a greedy error minimization algorithm. BayesianBits [34] proposes a quantization-based compression framework, which relies on residual error addition to power-of-two weights (including 0 for pruning). However, most of the above techniques either involve differentiable optimization, thus moving away from the post-training setting, or require fine-tuning/retraining to recover the DNN accuracy.

Our differentiators from the state of the art are many-fold: (i) we are the first to combine fine- and coarse-grained pruning techniques in our exploration phase. As we will demonstrate, the pruning granularity choice for a given DNN is not trivial and strongly depends on the model. (ii) We jointly prune and quantize each layer to a different sparsity ratio and precision. (iii) We directly include hardware feedback, from a DNN accelerator during the training of the RL agent. (iv) We do not employ any retraining or fine-tuning, during or post-exploration. A taxonomy of the above differentiating factors is presented in Table 1.

3 DNN COMPRESSION: PRUNING AND QUANTIZATION

In this work, we jointly apply pruning and quantization to a given DNN with fixed architecture, aiming to reduce its energy consumption, whilst maintaining high accuracy levels. Pruning directly compresses the model architecture, by removing parameters in a structured (i.e., coarse-grained) or sparse (i.e., fine-grained) manner. In contrast, quantization leaves the architecture intact and only reduces the precision used for the numerical representation of its parameters. Naturally, the two methods are non-overlapping and complement each other in reducing the resource budget of the targeted DNN. In fact, centroid-based quantization can directly benefit from a pruned model, as the quantization error can be reduced when less centroids are used [26]. However, both techniques exhibit a non-trivial trade-off between accuracy and energy consumption. Thus, their combined application requires a systematic approach (more details in Section 4). The rest of this section explains and motivates the intuition behind employing diverse pruning techniques and mixed-precision quantization in our framework.

3.1 Diverse Set of Pruning Techniques

Here, we motivate the need for employing diverse pruning algorithms, depending on the accuracy/energy trade-off of each given DNN. Typically, pruning algorithms follow either a fine-, or a coarse-grained approach, each offering unique hardware benefits. The intuition behind fine-grained pruning [4]–[6] derives from the observation that removing weights with small magnitude, e.g., close to zero, marginally affects the DNN accuracy. Considering weight-stationary inference, energy savings derive from the reduced switching activity (and consequently, lower dynamic power) of the arithmetic units of the accelerator, when operating on zero inputs [3]. Contrarily, coarse-grained pruning algorithms [7], [35] remove regular structures within convolutional layers (e.g., filters and channels), thus significantly reducing the model’s size and the number of operations. In this case, the entire computation is skipped, leading to easier exploitation of pruning and direct energy gains. Regarding compression trade-offs, weight pruning generally achieves much higher compression rate than structured pruning at comparable accuracy [20]. However, fine-grained sparse

TABLE 2: The pruning algorithms considered in our work.

Algorithm	Pruned Patterns
Sensitivity [5]	Weights
Level [4]	Weights
Splicing [6]	Weights
L1-Ranked [7]	Filters/Channels
L2-Ranked [7]	Filters/Channels
Bernoulli [36]	Filters
FM Reconstruction [35]	Channels

models cannot be supported by off-the-shelf libraries [9], and specialized hardware (e.g., sparse and zero skipping accelerators) and software solutions are needed to take full advantage of the sparsity and eventually improve inference.

Our framework is build upon the state-of-the-art pruning algorithms of Table 2. The corresponding pruned patterns of each technique are also presented. As these techniques prune convolutional or fully-connected layers, we focus on Convolutional Neural Networks (CNNs) in our work. Even though new pruning techniques can be seamlessly integrated, to extend the applicability of our framework to other architectures (e.g. RNNs), they are beyond the scope of this paper. Different DNNs respond differently to each technique and a single pruning method cannot provide an optimal solution across a range of models. Therefore, considering a diverse set of pruning techniques is crucial for the compression efficiency of our framework.

As a motivation example, Figure 1 compares the accuracy loss-energy savings trade-off for different sparsity levels (i.e., percentage of zero parameters) across varying DNNs. The pruning algorithms “Level” [4] (fine-grained) and “L1-Ranked” [7] (coarse-grained) are examined. An Eyeriss-based [2] DNN accelerator is considered and three different DNNs, trained on the CIFAR-10 dataset, are used: (a) VGG16, (b) ResNet50, and (c) MobileNetV2. As shown in Figure 1, all three DNNs have different sensitivity to pruning. MobileNetV2 is the most sensitive one, with L1-Ranked dropping accuracy more than 1.5% even for just 10% sparsity. Contrarily, ResNet50 proves the most resilient model, offering accuracy loss of less than 1% for almost all sparsity rates and both algorithms. Overall, the coarse-grained pruning algorithm exhibits significant accuracy loss (especially for rates above 40%) but also higher energy savings across all examined networks. However, this is not always the case. Assuming a conservative accuracy loss threshold of 1%, “Level” [4] for MobileNetV2 achieves 50% sparsity with higher energy gain compared with the *only* solution of “L1-Ranked” [7] (at 10% sparsity) that satisfies this threshold. Based on the aforementioned observations, the decision of a pruning technique is not trivial, it depends on the accuracy and energy requirements and is DNN-specific. Overall, relying on a single pruning algorithm will deliver suboptimal accuracy/energy trade-offs.

3.2 Mixed-Precision Quantization

Quantizing a (pruned) model introduces quantization errors to the (unpruned) parameters, causing a (further) detriment to the DNN accuracy. Mixed-precision quantization can

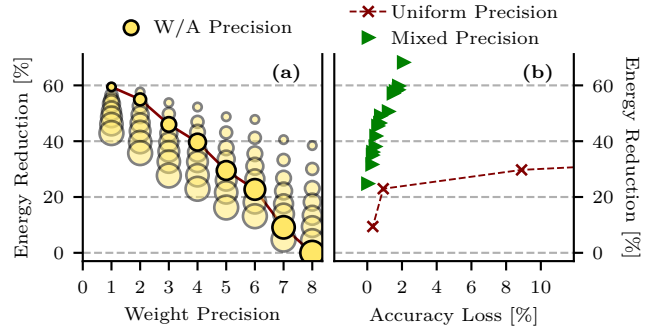


Fig. 2: (a) Energy reduction due to quantization when considering a fixed precision 8-bit Eyeriss-based [2] DNN accelerator, for different weight and activation precisions. x-axis represents the precision for the weights while the radius of each circle represents the precision of the activations. The solid line represents equal weight/activation precision. (b) Energy-Accuracy trade-off when considering Uniform vs Mixed per-layer precision for ResNet18 on CIFAR-10.

alleviate some of the associated accuracy loss (or allow for further potential gains), compared to the more common, uniform quantization. Note, mapping the entire model to the same precision may not be optimal, since not all layers have the same sensitivity to the introduced error [17].

To put the energy gains of mixed-precision quantization into perspective, we assume without any loss of generality a *worst-case scenario* in which the targeted DNN accelerator features a fixed architecture with single-precision MAC units, e.g., [1], [2], [37]. The precision of 8 bits is selected, since it is mainstream in embedded DNN accelerators (e.g., [37]), and it has been repeatedly demonstrated that 8-bit inference delivers close-to-floating-point inference accuracy [1]. As a result, in our worst-case scenario, computational energy gains due to mixed-precision quantization originate only from the precision-scaled inputs and the associated reduced toggling of the MAC units. Supplying 8-bit MAC units with lower-precision inputs (<8bit) leads to reduced switching activity and thus lower dynamic power [3]. Figure 2a presents the potential energy reduction from quantizing a DNN below 8 bits, on an 8-bit DNN accelerator. Energy gains are reported with respect to the baseline 8-bit Eyeriss-based [2] accelerator (details regarding the experimental setup and the energy evaluation are found in Section 5). As shown in Figure 2a, significant energy savings can be achieved when considering low-precision (<8-bit) weights and activations. For example, for 5-bit weights and activations, 29% energy reduction is obtained. Finally, Figure 2b presents an illustrative example that highlights the significance of considering mixed-precision compared to uniform quantization. As shown in Figure 2b, mixed-precision solutions populate a higher Pareto front, due to fine-grained manner in which the energy-accuracy space is explored. It is worth mentioning that for only 0.5% accuracy loss, mixed-precision quantization achieves 38.1% energy gains compared to only 9.4% for uniform quantization.

Inherently, transprecision DNN accelerator architectures, e.g., [38], [39], can potentially exploit the benefits originating from mixed-precision quantization in full, and maximize the energy gains. Note that although we consider a single-precision DNN accelerator, transprecision architectures can be seamlessly used in our framework due to its hardware-

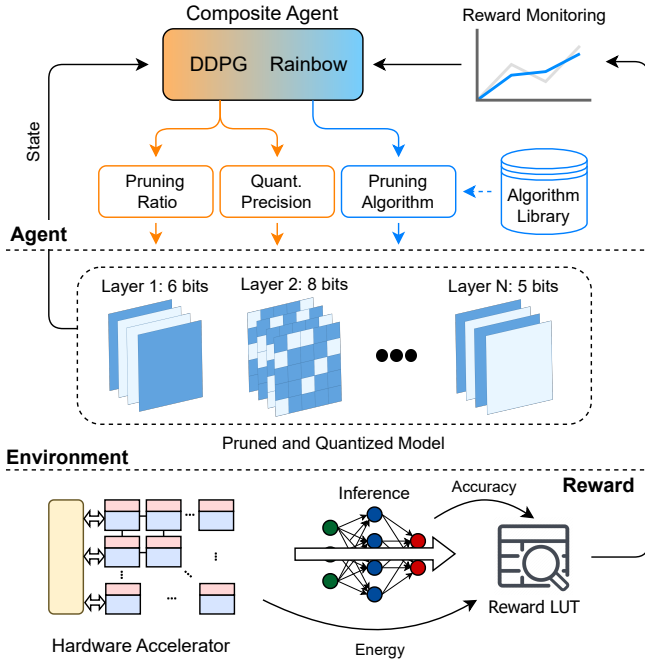


Fig. 3: Abstract overview of our proposed framework.

aware nature (see Section 4.2.3). In that case higher energy gains would also be expected.

4 PROPOSED HARDWARE-AWARE DNN COMPRESSION

In this section, the core of our proposed framework is detailed. We tackle the hardware-aware search for a joint pruning and quantization profile of a given pretrained DNN as an RL problem. Our goal is to learn, per DNN layer, the optimal (a) pruning ratio, (b) precision for weights and activations, and finally, (c) pruning technique.

An abstract overview of our framework is presented in Figure 3. Our composite agent receives the state of each layer and outputs three actions corresponding to its three compression directives (a-c). Accuracy and energy consumption are measured via validation inference and a custom energy model for the targeted hardware accelerator, respectively. Based on those measurements, we are able to deterministically formulate our reward in a LUT-based manner. Our agent strives to maximize the accumulated reward throughout a series of training episodes. Our framework outputs a DNN model, pruned with various techniques and quantized with mixed precision, which features a good trade-off between energy consumption and accuracy.

In contrast to the state-of-the-art approaches, our framework avoids any retraining or short-term fine-tuning. This can potentially enable on-device optimization and DNN compression, as opposed to the computationally intensive alternative that requires fine-tuning. For example, on-device training of the required RL-agents (see Section 4.2) is significantly less computationally intensive than (re)training a DNN model (e.g., ResNet50 on ImageNet). In addition, on-device optimization preserves data locality and therefore satisfies user privacy constraints. Finally, the applicability of

our framework is extended to scenarios where training data are either unavailable or proprietary.

The following subsections contain detailed descriptions of each aspect of our framework, as portrayed in Fig. 3.

4.1 Layer-wise Pruning and Quantization

Regarding our one-shot pruning and quantization technique, several considerations are detailed below. Utilizing a combination of fine- and coarse-grained pruning techniques (see Table 2) comes with the additional challenge of resolving dependencies on more complex architectures (e.g. ResNets). For example, pruning the last layer of a residual block with a structured technique enforces an identical pruning action at the shortcut layer of the same block. Instead of simply avoiding pruning such layers, we have the flexibility to sparsify them using fine-grained techniques, thus further reducing the overall energy consumption without causing structural mismatches. This reveals the importance of conducting layer-wise pruning (i.e., dependencies are resolved at the first dependent layer). All layers are subject to pruning (and quantization) in our framework.

Considering that no retraining is conducted in our framework, the accuracy of the compressed DNNs will not improve post-optimization. Thus, a conservative pruning policy is expected to maintain acceptable accuracy levels, especially in the coarse-grained case. For this reason, quantization is applied upon the pruned model, as a second step, to further reduce the energy consumption. In our framework, we apply per-channel, asymmetric, linear, post-training quantization with activation clipping based on a Laplace distribution [21]. We set the same precision for both weights and activations, in order to reduce the vast size of our joint design space. Note that, quantization is always applied to at least the default precision of the targeted accelerator, i.e., 8-bit in our case.

4.2 Automated Compression with Composite RL Agent

To effectively traverse the large design space of possible pruning and quantization configurations, we propose a novel composite RL agent, able to learn the optimal compression profile for each DNN layer. Our agent consists of a DDPG [40] and a Rainbow [41] agent, each assigned to a different task: the DDPG algorithm is responsible for learning the pruning ratio and quantization precision, whereas Rainbow controls the pruning algorithm. Both agent components are equipped with a prioritized replay buffer, to favor experiences with higher temporal difference (TD) error, thus improving the sampling efficiency of our framework. A schematic overview of our RL scheme is demonstrated in Figure 4. In the following subsections, we describe in detail how each part of our composite agent is trained to achieve its allocated task.

4.2.1 DDPG agent

DDPG [40] is an off-policy algorithm, equipped with an actor-critic model. The actor is responsible for delivering the agent's actions (i.e., pruning ratio and quantization precision), whereas the critic is a Q-value network that determines the quality of the actor's actions, based on the gathered reward. We enhance the learning procedure with

target critic and actor networks, and the use of truncated normal distribution for noise exploration purposes. The DDPG algorithm operates over a continuous, multidimensional action space, which in our case is a two-dimensional vector within the range of $[0, 1]$ that corresponds to the pruning ratio and selected precision (i.e., the outputs of the actor). Using a continuous space allows for finer control over both variables. In order to translate the continuous actions to precision (i.e., bits), a simple linear mapping is required, followed by rounding to the nearest integer. The environment state, which serves as the observation for the agent, is a 13-dimensional vector s_t and represents the full embedding of layer t . In other words, the state vector mirrors every characteristic of each layer, so that the actor can act based on a rich, fully-descriptive set of features. In the case of a convolutional layer:

$$s_t = \{t, 0, C_o, C_{in}, h_{in}, w_{in}, str, k, E_t, P_t, M_t, E_t^{red}, a_{t-1}\}, \quad (1)$$

where t is the layer index (equivalent to the agent step), $C_{in} \times h_{in} \times w_{in}$ is the size of the input feature maps (FM), C_o are the output FM channels, k is the kernel size, str is the stride, E_t is the energy consumption of (unpruned and non-quantized) layer t , P_t is the number of its weight parameters (linked with pruning effects), M_t is its memory size (i.e., $P_t \times q_b$, where $q_b = 32$ precision bits for floating point weights, linked with quantization effects), E_t^{red} is the reduced energy of layer t caused by actions a_t and a_{t-1} are the actions for layer $t - 1$ (i.e., at step $t - 1$). Similarly to (1), for a fully connected layer:

$$s_t = \{t, 1, M, N, h_{in}, w_{in}, 0, 1, E_t, P_t, E_t^{red}, a_{t-1}\}, \quad (2)$$

where $N \times M$ is the shape of the weight tensor.

4.2.2 Rainbow agent

We implement a Rainbow agent [41] to learn the optimal pruning technique for each layer of the input DNN. Although, both the pruning ratio and quantization precision can be easily represented as continuous variables, pruning algorithms can only be interpreted as an array of indexes. Incrementing or decrementing indexes does not convey any physical information (importance) and is therefore hard to optimize with the continuous action space of DDPG. For this purpose, we leverage the inherently discrete action space of the Rainbow algorithm. Rainbow improves upon Double Q-Learning with a dueling model architecture [42], i.e., a value and advantage subnetworks, which introduce attention to the selected actions. The injection of noise in the last layer of both subnetworks enriches the agent's robustness to perturbed observations from its environment. Finally, instead of using the absolute values of the outputs from the dueling architecture, an estimation of their distribution is obtained, which boosts the agent's overall efficiency.

Rainbow and DDPG are directly connected, forming our composite RL agent. The input (i.e., state vector) to the Rainbow model is the the output of the feature extractor of our compression policy, i.e. the output of the hidden layer of the DDPG actor network (see Figure 4). At each step, after DDPG is updated, the outputs from the hidden layer of its actor network feed the Rainbow model as inputs. These

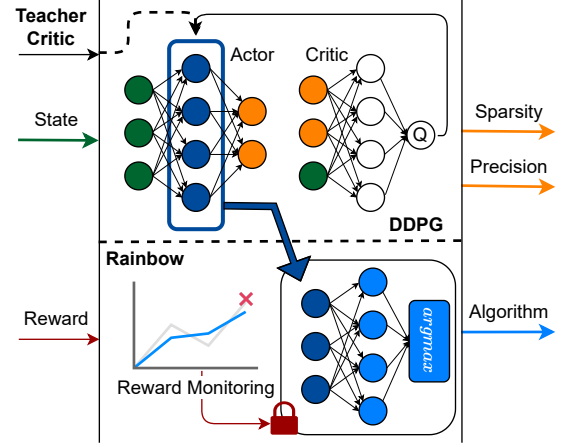


Fig. 4: Overview of our composite RL scheme, including our two agents (DDPG and Rainbow). The reward monitoring component, which activates the Rainbow agent, and the teacher critic model which enriches the DDPG update, are also depicted.

inputs/features are used to produce the action: a discrete pruning technique. Then, the Rainbow subnetworks (value and advantage) are allowed to update. This feedback significantly helps the Rainbow model in learning abstract features from the DNN, without any computational overhead.

Importantly, the DDPG needs to (start to) learn the pruning sensitivity of each layer *before* Rainbow appoints the optimal pruning algorithm. Our aim is for Rainbow to safely extract patterns from its input, which would be suboptimal if the input features were rapidly changing (i.e., at the initial phase of DDPG training). Thus, we freeze the training process of Rainbow, until the DDPG feature extractor has shown signs of improvement (i.e., increased moving average reward). Random pruning techniques (from Table 2) are sampled at this point to remove any bias towards any specific technique. We implement a light-weight reward monitoring scheme, which keeps Rainbow stagnant throughout the primary exploratory period of our optimization. As soon as the reward/episode curve reflects a consistent improvement, the Rainbow agent is unlocked and allowed to control the selection of pruning techniques using the already mature features of the DDPG actor backbone. The loss of the Rainbow agent does not back-propagate, to avoid any bias overestimation pitfalls. Note, rewards are fed to the agent at every step, since the Rainbow model requires an update before each compression action. Overall, the Rainbow agent learns to associate abstract features of pruning and quantization (i.e., from the last layer of the DDPG actor's feature extractor) with the best fitted technique for these features.

4.2.3 Hardware-aware reward

We continue with the description of our reward mechanism, shared between both parts of our composite agent (i.e., DDPG and Rainbow). We formulate our reward as a Look-up Table (LUT), indexed by accuracy and energy consumption measurements. Accuracy is evaluated by running inference on a validation subset. Simultaneously, we estimate the energy consumption of the pruned and quantized DNN, using a custom energy model (see Section 4.3). Energy mea-

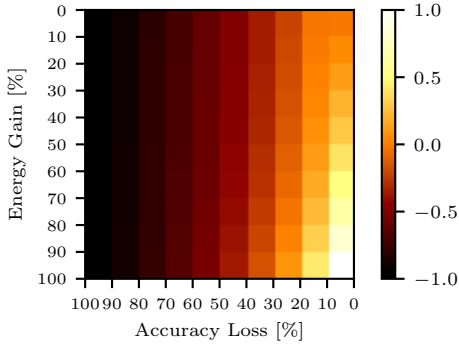


Fig. 5: Heatmap of our LUT-based reward. x-axis contains the accuracy loss and y-axis the energy gain. The reward is significantly higher for solutions of accuracy loss $< 10\%$, which is a realistic target region of our framework. For readability, the heatmap is plotted at 25% of its actual resolution.

measurements serve as direct feedback from the environment to the agent, as part of the reward. Note that, although we target energy efficiency in our work, any other hardware metric (e.g., latency) is seamlessly supported since it can be measured in an identical manner. Naively formulating our reward as an analytical notation comprising the accuracy and energy consumption (e.g., a linear combination of the two measurements), could lead to DNNs of significant accuracy loss, albeit high energy gains. Since our framework does not retrain the final compressed model, to improve its prediction capabilities, we are realistically interested in a limited, high-accuracy region in the Pareto curve of the possible accuracy-energy trade-offs. Thus, we create a LUT of size 40×40 , containing a satisfactory number of accuracy loss and energy gain combinations (w.r.t. the dense model), heavily favoring, intentionally, solutions with small accuracy loss. A heatmap of our custom reward function is illustrated in Figure 5. Note, a sub-sampled version of our reward function is depicted, at 25% of its actual resolution, for readability purposes. Our LUT-based reward is significantly higher when the accuracy loss is less than 10%. Hence, the RL agents are strongly incentivized to prioritize actions that lead to minimal accuracy loss. Similarly, the reward corresponding to minimal energy gains ($< 5\%$) but also small accuracy loss ($< 5\%$), is a small negative number, to slightly discourage RL agents from outputting close-to-zero compression actions.

4.3 Energy Model

Here, we describe our custom energy model to estimate the energy consumption of the targeted accelerator platform and for a given pruned and quantized DNN workload. We operate on a layer-wise fashion to extract an energy measurement for each DNN layer. Each measurement comprises two parts: the energy consumption related to memory accesses and data movement (E_{mem}) and the one related to MAC computations (E_{comp}). Similar to [13], E_{comp} is approximated as the number of MAC operations weighed by the energy cost of a single computation. E_{mem} on the other hand is calculated by extracting the total number of accesses to the accelerator’s memory, multiplied by the energy cost of a single memory access. We enhance this model by applying a multiplicative factor to both terms,

which conveys the necessary information about the effects of pruning and quantization to the specific layer. These factors (R_x) are appropriately named reduction coefficients, as their purpose is to apply a reduction to the energy consumption due to compression. Thus, the total energy (E_{total}) of a DNN model with L layers is given as follows:

$$E_{total} = \sum_{l=0}^L E^l = \sum_{l=0}^L (E_{mem}^l + E_{comp}^l), \quad (3)$$

$$E_{mem} = \#acc * e_{mem} * R_{mem}, \quad (4)$$

$$E_{comp} = \#comp * e_{comp} * (R_{pruned} + R_{unpruned}), \quad (5)$$

where e_{mem} and e_{comp} are the cost values for a single memory access and MAC computation, respectively; R_{mem} , R_{pruned} and $R_{unpruned}$ are the reduction coefficients for the memory accesses, and for computations with pruned and unpruned parameters, respectively. Note, the coefficients are non-unit and valued in the closed set of $[0, 1]$, such that an energy reduction w.r.t the baseline is guaranteed. Below, we analyze the calculation of each term in the equations above.

Both the number of MAC operations ($\#comp$) and memory accesses ($\#acc$) are automatically provided by an open source tool¹, namely NN-Dataflow. Inspired by works related to efficient mapping of DNNs to hardware platforms [43], [44], NN-Dataflow offers a fast exploration mechanism upon diverse dataflow scheduling choices including array mapping, loop blocking and reordering, and (coarse-grained) parallel processing within and across layers. Noticeably, the dataflow optimization techniques for intra-layer parallelism and inter-layer pipelining presented in Tangram [44] are incorporated. Such a direct mapping and measurement helps in capturing all unique features of the accelerator, leading to an accurate modeling of energy-aware inference. Finally, the single cost values in (4),(5) can be estimated by measuring the cost of running a single MAC (floating-point) operation on the accelerator, as well as the energy consumed by one access to memory.

Reduction coefficients incorporate the effect of pruning and quantization on each comprising term of (3), and their values depend on the selected pruning algorithm and precision. Quantization is capable of reducing the computational cost of MAC operation, due to the lower precision, for all unpruned parameters. To calculate that effect, we evaluate the energy consumption of the 8-bit MAC units when operating over reduced-precision multiplicands (i.e., < 8 -bit weights and activations). To that end, we design and synthesize an 8-bit MAC circuit that consists of an 8-bit multiplier and a 32-bit adder to avoid accumulation overflow. Gate-level timing simulations with realistic inputs obtained from quantized networks at different precisions are performed. All the possible combinations for the precision of the multiplicands (i.e., ≤ 8 bit) are examined, and for each combination we run a separate simulation to obtain the respective switching activity of the MAC unit. Finally, the obtained switching activities are used to measure the power consumption of the MAC unit for each precision combination. In this way, we are able to derive the computational power reduction ratio R_Q between the baseline (i.e., 8-bit precision for weights and activations) and any

1. github.com/stanford-mast/nn_dataflow

given quantization codebook (namely, Q_W and Q_A bits for weights and activations, respectively):

$$R_Q = \frac{P_{Q_W/Q_A}}{P_{8/8}}, R_Q \in [0, 1]. \quad (6)$$

Pruning effects to energy consumption differ according to the selected pruning algorithm. Fine-grained pruning allows for moderate reduction in the computation-related energy, whilst demanding the same amount of memory accesses as the baseline, since parameters set to zero still need to be fetched from memory to calculate intermediate results for sparse matrix arithmetic operations. Additionally, computations with pruned parameters (i.e., set to 0) have a non-zero energy cost. The intuition behind this statement is that when multiplying by 0, the power consumption (and thus energy) of the MAC unit is expected to be significantly reduced but it will not be zero, due to the static power consumption as well as any toggling activity of the 32-bit adder. As mentioned in Section 3, specialized hardware (e.g., sparse and zero skipping accelerators) and software solutions are needed to take full advantage of the induced sparsity by fine-grained pruning. To calculate this non-zero quantity, we evaluate the energy consumption of the MAC unit when one of the multiplicands is 0 (e.g., as in fine-pruning), similarly to the aforementioned power estimation technique for quantization. In this case, the other multiplicand and the partial sum are free to toggle. Thus, we come up with a “penalty” term P_{FG} which represents the average energy consumption of computations with pruned parameters compared to unpruned ones, when using a fine-grained pruning algorithm. According to our calculations, this value is set at 0.2, i.e., the energy used in MAC operations with weights set to 0 is reduced by 80% compared to computations with unpruned parameters. Putting everything together, in the case of fine-grained pruning, the coefficients from (3),(4),(5) take the following form:

$$\begin{aligned} R_{mem} &= 1, \\ R_{pruned} &= P_{FG} * S, \\ R_{unpruned} &= (1 - S) * R_Q, \end{aligned} \quad (7)$$

where S symbolizes the fine-grained sparsity ($S \in [0, 1]$). (7) shows that fine-grained algorithms may lead to moderate energy gains, due to the “penalized” computational energy consumed of unpruned parameters.

Coarse-grained pruning algorithms provide a straightforward reduction in both the necessary MAC computations and accesses to memory, directly proportional to the induced sparsity. Their structured nature allows for entire blocks (i.e., filters/channels) to be removed and consequently for parts of matrix multiplication operations to be skipped. The corresponding patterns in data movement are analogously simplified. Summarizing, such pruning algorithms correspond to the following reduction coefficients:

$$\begin{aligned} R_{mem} &= 1 - S, \\ R_{pruned} &= 0, \\ R_{unpruned} &= (1 - S) * R_Q. \end{aligned} \quad (8)$$

The above mathematical formulation of our energy model allows for quick energy estimations during RL training.




Fig. 6: Schematic overview of the Eyeriss-based [2] targeted DNN accelerator for our evaluation.

5 RESULTS AND ANALYSIS

5.1 Experimental Setup

In this section we assess the efficiency of our proposed framework by performing an in-depth evaluation over nine state-of-the-art DNNs and three image classification datasets. For our analysis we consider: VGG11, VGG13, and ResNet18 trained on CIFAR-10, VGG16, ResNet34, and MobileNetV2 trained on CIFAR-100, and VGG19, ResNet50, and SqueezeNet trained on ImageNet. Our RL agents are trained for 1100 episodes (first 100 constitute the warm-up), sampling 64 experiences per update from the prioritized replay buffer, whose size is set at 1000 experiences. Both actor and critic DDPG networks are built with 3 hidden fully-connected layers of 300 neurons and are trained with learning rates of 10^{-3} and 10^{-4} , respectively. Noise for the truncated normal distribution is initialized at 0.6 and after the warm-up is decayed with a factor of 0.99 after each episode. A discount factor of 1 is used. The (numerous) hyperparameters required to configure the Rainbow part of our agent are directly taken from [41].

During our optimization, rewards are computed using a 10% subset of the validation set, regarding the accuracy component of the LUT. As a target hardware platform, we consider an 8-bit Eyeriss-based DNN accelerator [2], which is a tile-based architecture (see Fig. 6). Specifically, we consider a 64×64 2D array of processing elements (PEs) for each tile, with a local register file of 64 bytes in each PE, and a shared global SRAM buffer of 32KB, similar to [44]. A standard 2D hierarchical memory architecture is utilized (i.e., memory channels are assumed to be on the four corners of the chip) and a total DRAM bandwidth of 3.2Gbps. Energy measurements from the hardware accelerator are extracted using the custom energy model described in Section 4.3. The MAC units are designed in Verilog-RTL, synthesized with zero-slack, and mapped to the 7nm ASAP7 library. The optimized arithmetic components of the commercial Synopsys DesignWare library are used for the adder and multiplier. Synopsys Design Compiler and the `compile_ultra` command are used for circuit synthesis, QuestaSim is used for gate-level timing simulations (and obtaining the respective switching activities), and Synopsys Primitime for measuring the power consumption. Finally,

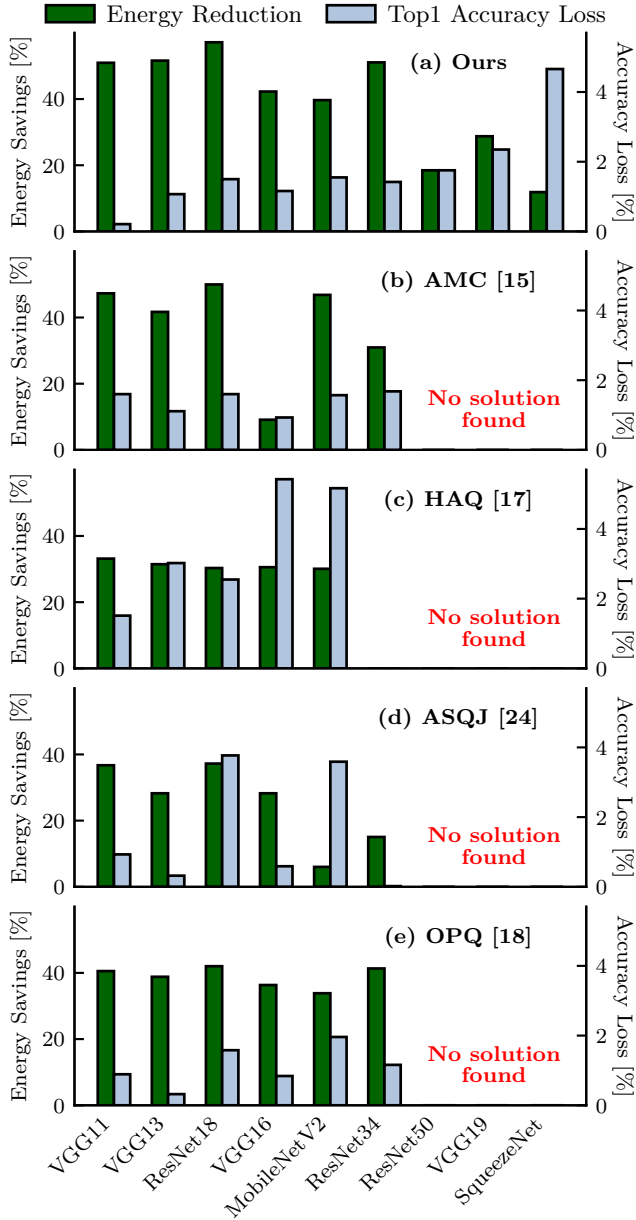


Fig. 7: Energy gain vs accuracy loss evaluation for (a) our framework, (b) AMC [15], (c) HAQ [17], (d) ASQJ [24] and (e) OPQ [18]. Without retraining (or with only a single epoch for OPQ), the state of the art technique could not converge to a compressed DNN of acceptable accuracy loss ($< 10\%$) for the ImageNet dataset.

the top-1 classification accuracy of the obtained compressed DNN is evaluated on the test set.

5.2 Accuracy/Energy Trade-off

To evaluate the efficiency of our framework and put the delivered energy-accuracy trade-off into perspective, we compare the DNNs produced by our framework against the following state-of-the-art frameworks: (i) *AMC* [15] trains a DDPG agent to apply per-layer channel pruning, (ii) *HAQ* [17] applies mixed-precision quantization to weights and activations, also utilizing the DDPG algorithm, (iii) *ASQJ* [24] jointly prunes and quantizes DNN weights, using the ADMM technique, and (iv) *OPQ* [18] applies one-shot pruning and quantization using an analytical

model on pretrained DNN weights. We select the aforementioned approaches to cover a wide spectrum of comparisons: AMC and HAQ serve as standalone approaches, applying solely pruning and quantization, respectively, in a learning-based manner. They also constitute hardware-aware approaches (see Table 1). Contrarily, both ASQJ and OPQ tackle the joint pruning/quantization design space, the former using ADMM and the latter with an analytical model. Additionally, the first three approaches attempt to learn the compression profile via parameter update (i.e., RL or ADMM), whereas OPQ approximates the error minimization function via the Lagrange multiplier. Finally, OPQ is an one-shot method, i.e., no access to the training data is required to calculate the pruning masks or the quantization codebook. All accuracy loss and energy reduction measurements below are reported w.r.t. the baseline DNN (i.e., dense DNN quantized at 8 bits). Since AMC [15] uses floating-point inference, we quantize the resulting pruned DNN to 8 bits. The same quantization algorithm that is used in our framework is also used to quantize the baseline and AMC.

To the best of our knowledge our framework is the first learning-based and hardware-aware one that does not require fine-tuning or retraining. To be clear, one-shot compression via joint pruning and quantization has been proposed in prior work [18], [33], [34]. Even outside the one-shot environment, fine-tuning steps (as well as long-term retraining) are often used by related works [15], [17], [24] in the exploration phase. Retraining is orthogonal with any compression framework and acts as a standalone improvement technique to significantly recover accuracy loss without also affecting the energy efficiency of compressed model. Thus, to ensure a fair comparison, we do not conduct long-term retraining on any compressed DNN. Nevertheless, we allow for fine-tuning to be applied between the exploration steps of AMC [15], HAQ [17] and ASQJ [24]. For OPQ [18], a few fine-tuning steps are used post-compression. Although this gives a significant advantage to the state of the art, it enables us to implement the respective algorithms without tampering with their optimization flow.

Figure 7 presents our evaluation in terms of the accuracy loss vs energy gain trade-off of the compressed DNNs obtained from our framework, as well as the related works. The derived trade-offs and comparisons are analyzed in the following subsections.

5.2.1 Evaluation of our Framework

As shown in Figure 7(a), our framework delivers high energy savings and small accuracy loss across all DNNs and datasets. A conservative threshold of $< 2\%$ top-1 accuracy loss was imposed on our compressed DNNs for the CIFAR-10 and CIFAR-100 datasets. This is relaxed to $< 5\%$ for the more challenging classification task on the ImageNet dataset. On average, our framework achieves 53.23%, 44.33%, and 19.7% energy reduction for the CIFAR-10, CIFAR-100, and ImageNet DNNs, respectively. The corresponding average top-1 accuracy loss is 0.93%, 1.38% and 2.92%, on average. As expected, the compression-induced reduction on energy consumption is at its peak for the easier classification task, on the CIFAR-10 dataset. Error tolerance of these models stems from the fact that a substantial amount of parameters may be redundant, or quan-

tization errors can be masked from the last layer’s softmax probabilities. Our agent exploits that fact, as the accuracy loss is incorporated in its reward, and strives to compress such models more aggressively, driving the energy gains to high levels. On a similar but lower scale, the same can be said for the CIFAR-100 dataset. Interestingly, the lowest obtained energy gains are observed for the MobileNetV2 architecture, which includes lightweight convolution operations (e.g., depth-wise convolution). Such operations are constrained by design, and do not make good candidates for compressing their parameters. We expect that a DNN architecture of moderate depth but convolution operations of high volume would be ideal for exploiting the benefits of our framework to their maximum. A great example is the ResNet34 architecture, which even though is trained on the CIFAR-100 dataset, achieves an impressive 51.07% energy gain with only 1.4% top-1 accuracy loss.

Understandably, as the dataset/task becomes more complex/challenging (and since we do not use fine-tune/retraining) the obtained energy savings diminish. Though, it is noteworthy that even for ImageNet, where energy gains are limited without retraining, our framework still manages to produce models of considerable energy savings within the strict accuracy bounds ($< 5\%$). For example, for ResNet50 on ImageNet, we achieve 18.47% energy reduction and only 1.75% accuracy loss. Similarly, for VGG19 on ImageNet, the energy savings increase to 28.75% while the accuracy loss is 2.35%. Even in the case of SqueezeNet, which is a compressed architecture by design, we achieve 11.88% energy gain for 4.66% accuracy loss. Again, our framework delivers such energy gains by efficiently searching the joint pruning-quantization space and without requiring fine-tuning or retraining.

5.2.2 Comparison against Standalone Approaches

In this section, we compare our framework against the standalone pruning/quantization state-of-the-art methodologies: AMC [15] and HAQ [17], in Figure 7(b) and (c), respectively. As mentioned above, both approaches involve RL-based optimization for either the per-layer pruning ration (AMC) or precision for weights and activations (HAQ), coupled with hardware-aware reward. AMC reaches up to 49.9% energy gain for ResNet18 on CIFAR-10 with 1.6% accuracy loss. On average, AMC achieves 46.32% and 28.97% energy reduction, respectively, for the CIFAR-10 and CIFAR-100 DNNs, while the corresponding accuracy loss is 1.44% and 1.39%. In comparison, our work achieves 0.51% and 0.01% lower top-1 loss with 6.91% and 15.36% higher energy gain, on average. Quantized DNNs by HAQ feature 31.64% and 30.32% average energy reduction and average top-1 accuracy loss of 2.36% and 5.3%, for CIFAR-10 and CIFAR-100. Only a single use-case (VGG11 on CIFAR-10) adheres to the 2% accuracy loss threshold set by our framework. Interestingly, HAQ achieves semi-constant ($\sim 30\%$) energy gains over all the examined scenarios, but with the highest accuracy loss across all techniques. This highlights HAQ’s difficulty in detecting the sensitivity of each model to quantization, and its reliance on post-optimization retraining to reduce the accuracy loss below more conservative thresholds (e.g., 1%). Note, HAQ could not converge to a solution of acceptable accuracy drop ($< 10\%$) for ResNet34

on CIFAR-100. Similar to our framework, as the dataset complexity increases, the energy savings also decrease.

Our framework, outperforms the two related methodologies for all the examined DNNs except for AMC on MobileNetV2. This inherently compressed architecture is amenable to structured pruning due to the regular structure of depthwise separable convolutions. In contrast, weight pruning is considerably less impactful, as only a few redundant parameters exist. Hence, only conservative pruning can be applied without greatly hurting prediction accuracy, which leads to reduced energy efficiency. As a result, our framework achieves 40% energy reduction while AMC achieves 47% energy savings. However, this difference is specific to the structure of MobileNetV2 and subject to the fine-tuning used by AMC to recover the accuracy loss, enabling AMC to further boost the applied pruning.

It must be pointed out that neither AMC nor HAQ were able to identify a solution with less than 10% accuracy loss for ImageNet. Its challenging nature and the high complexity of the associated DNNs, make it significantly more difficult to apply pruning or quantization and still maintain the desired accuracy, without long-term retraining. This further highlights the significance of all the aspects of our work, i.e., considering many pruning algorithms and intelligently selecting the optimal one, applying diverse pruning and mixed-precision quantization in a hardware-aware manner and guiding the RL exploration to high accuracy solutions via our flexible LUT-based reward.

5.2.3 Comparison against Joint Approaches

Finally, we compare against the joint pruning/quantization state-of-the-art approaches. Figures 7(d) and (e) present the top-1 accuracy loss and energy gain of ASQJ [24] and OPQ [18], respectively. Note, ASQJ conducts an ADMM-based optimization where each compression stage is represented by a different variable update, whereas OPQ derives the pruning masks and quantization steps in one-shot using a Lagrangian-based analytical model. To have a fair comparison and achieve competitive results, we include 5 retraining epochs for OPQ, after the one-shot calculations, for reasons mentioned in Section 5.2. For ImageNet, this was reduced to 1 epoch.

On average, ASQJ achieves 34.07% and 16.45% energy savings and top-1 accuracy loss of 1.66% and 1.4%, for the CIFAR-10 and CIFAR-100 datasets, respectively. Notice, that the energy gains obtained by ASQJ are smaller compared to both AMC and our framework. This is expected, as fine-grained pruning (employed by ASQJ) reduces energy in a less impactful way than coarse-grained pruning (used in both our framework and AMC). Pruned and quantized DNNs from OPQ deliver on average 40.46% and 37.17% with only 0.93% and 1.32% top-1 accuracy loss. Such accuracy results are almost identical to the ones from our framework (see Figure 7(a)). Still, our energy gains for both datasets are on average 12.77% and 7.16% higher, respectively, thus proving the effectiveness of one-shot methodologies in discovering an optimal pruning and quantization profile per layer. Note however, OPQ requires 5 retraining epochs to reach such accuracy levels.

Once again, the state-of-the-art did not manage to produce a DNN of acceptable accuracy levels ($< 10\%$) on

the challenging ImageNet dataset. For ASQJ, the ADMM algorithm, despite its known capabilities, cannot traverse the design space as effectively as our RL agent, since it lacks the incentive to explore the limited region of high-accuracy solutions, which we achieve via our specialized reward function (see Section 4.2.3). OPQ on the other hand, even though no access to training data is needed, seems to heavily rely on fine-tuning steps, since the trend on CIFAR datasets did not continue here as well. Contrarily, our framework operates in a truly one-shot fashion, without needing any retraining, albeit at the cost of elevated time complexity (see Section 5.3.3).

5.3 Exploration Efficacy

5.3.1 Pruning/Quantization Analysis

In this section, we provide insight to the intuition behind our RL agent’s actions, regarding the pruning ratio, quantization precision and pruning algorithm, per layer. Figure 8 presents a descriptive example of the solutions obtained by our framework and the derived pruning and quantization policy. For readability reasons, ResNet18 is used (comprises only 20 layers), but similar results are obtained for the other DNNs. Interpreting the selected pruning policy, the first layers are pruned conservatively using coarse-grained algorithms (see Table 2). Considering the sensitivity (and therefore, lack of redundancy) of the first layers, a conservative pruning policy is a reasonable outcome, in order to retain the prediction accuracy to acceptable levels. We observe that a series of coarse-grained algorithms (at similar pruning ratios) are selected for the middle stage of the model, since the size of each layer remains the same within that stage. Additionally, our Rainbow agent recognizes the DDPG features corresponding to low sparsity and attributes coarse-grained pruning techniques to increase the potential energy gains. Interestingly, the last two layers, which are fully-connected, and consequently more redundant, are pruned in a more aggressive manner with fine-grained techniques. Our quantization policy has protected the first and last layers from aggressive quantization, keeping them at 7 or 8 bits. Noticeably, the final shortcut convolutional layer (i.e., layer 16 in Figure 8) is barely pruned, but heavily quantized to low precision. Extensively pruning that layer would result to accuracy degradation, due to its inter-layer dependencies to the consecutive residual block. Thus, quantization drives the potential energy gains without significantly deteriorating the accuracy by pruning a critical layer. The heterogeneity of sparsities/precisions directly reflects our agent’s ability to identify the sensitivity in pruning and quantization of the different layers of the DNN. Moreover, the complementary nature by which each layer is pruned and quantized stands as testament to the efficacy of our framework in intelligently exploring the joint design space.

5.3.2 Comparison against heuristic algorithm

To better highlight the performance of our learning-based design space exploration, we compare our composite RL scheme against the widely adopted NSGA-II [45] algorithm. NSGA-II has prevailed as one of the most useful and fast heuristic-based exploration algorithms in the literature,

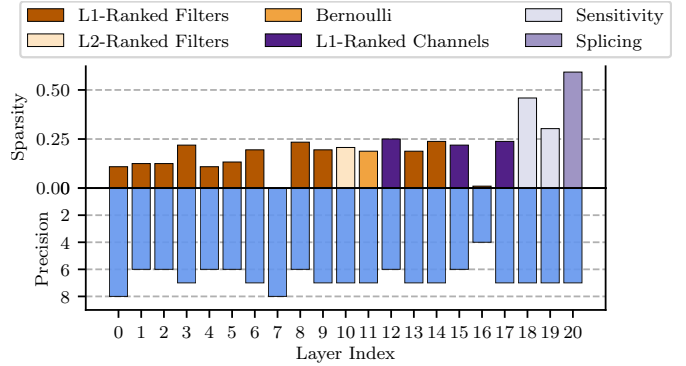


Fig. 8: Analytical demonstration of the pruning and quantization decisions made by our RL agent for the ResNet18 model on the CIFAR-10 dataset.

due to its capability of converging quickly and generating Pareto-optimal solutions.

NSGA-II searches for an optimal pruning/quantization configuration simultaneously for all L DNN layers (contrary to the step-wise operation of RL algorithms), thus creating a genome of length $3 \times L$. Continuous variables are therefore required, forcing the pruning algorithm variable to be controlled by rounding to discrete values which represent indices. All (initially random) chromosomes (i.e., parent population) are subjected to the standard iterative procedure of the NSGA-II, i.e., tournament selection, simulated binary crossover and polynomial mutation. Such operators are tailored for continuous space exploration. From a combined pool of parent and children chromosomes, the non-dominated solutions are selected via fast non-dominated sorting and truncation based on crowding distance. To ensure a fair comparison, the GA is allowed the same number of evaluations (i.e., energy and accuracy estimation) as our RL algorithm. Since our agents are trained for 1100 episodes (see Section 5.1), we configured the NSGA-II to run for 55 generations, with a population size of 20 chromosomes, to balance the exploration/exploitation trade-off. Importantly, all the evaluations are conducted in the same manner as our technique, including our hardware aware reward, which serves as the fitness function² (i.e., a single fitness objective is conducted). The output of the genetic algorithm is the single solution (pruned/quantized DNN) with the highest reward value (see Section 4.2.3).

We present the comparative evaluation results in Figure 9. Overall, NSGA-II produces DNNs of high accuracy loss, albeit elevated energy gains. For CIFAR-10, which is a more error-tolerant task (classification to 10 classes), a significant (but not destructive) top-1 loss of $\sim 8\%$ on average is observed. For the rest of the extracted models, for the more demanding datasets (with faded markers in Figure 9), the best NSGA-II solution fails to limit the classification accuracy below 10%. This can be attributed to the limited amount of allowed evaluations and the imposed tight accuracy constraints. As mentioned in Section 4.2.3, the area of interest for our reward is a very limited, high-accuracy region in the Pareto curve of the possible accuracy-energy trade-offs. That is because training data are not used

² Since NSGA-II strives to minimize the fitness objectives, we use the inverse reward (i.e., multiplied by -1).




Fig. 9: Comparative evaluation between our composite RL agent and the NSGA-II algorithm [45], in terms of top-1 accuracy loss and energy gain. NSGA-II does not manage to produce DNNs of acceptable accuracy loss ($< 5\%$), and the CIFAR-10 networks are the only ones with loss lower than 10%. Faded lines and markers represent DNNs of unacceptable accuracy loss ($> 10\%$).

and therefore, the possibility of improving the final accuracy after pruning or quantization is absent. Our learning-based exploration is able to guide the RL agent towards (useful) solutions of high accuracy (which represent a very small sample of our entire design space), even at the cost of moderate energy gains (e.g., for the ImageNet dataset). Genetic algorithms are less sample efficient (since they heavily rely on stochastic processes) in favor of quick execution, and cannot respond well when having to deal with a limited amount of evaluations.

5.3.3 Execution Time and Memory Requirements

Finally, we evaluate the execution time and memory requirements of our framework, and specifically, our RL-based optimization, in comparison to state-of-the-art techniques [15], [17], [18], [24]. Since the number of iterations (e.g., RL training episodes) is a user-defined parameter, our comparisons involve a single iteration for each studied method. To ensure a fair comparison against the state of the art, the same contributing factors to the performance of each technique were used (e.g., number of worker threads, batch size, etc.). Measurements were conducted on a desktop computer featuring an NVIDIA GeForce RTX 2080 Super GPU operating at 1.65 GHz and 32GB of RAM. Table 3 presents the normalized (w.r.t. the fastest technique) average execution time of a single iteration for all evaluated models on each dataset, taken from an average over 50 iterations, to avoid statistical bias. Table 4 presents a comparison of the memory utilization for the studied techniques.

Overall, the reported execution time of our framework is on the higher end compared to the state of the art, whilst having comparable memory requirements. OPQ [18] presents the lowest time complexity due to its quick one-shot mathematical analysis of pruning/quantization, which translates though to elevated memory utilization for ImageNet, mostly due to multiple copies of weight tensors. ASQJ’s [24] memory efficiency due to the sequential nature of parameter update of the ADMM algorithm proves to be

TABLE 3: Comparative evaluation for the execution time of a single iteration among our proposed framework and state-of-the-art techniques. The average execution time over all evaluated models on each dataset is obtained. The fastest technique is used as the normalization baseline for each dataset.

Dataset	Ours	AMC [15]	HAQ [17]	ASQJ [24]	OPQ [18]
CIFAR-10	5.61x	5.60x	2.743x	19.89x	1.00x
CIFAR-100	29.35x	12.85x	4.94x	19.82x	1.00x
ImageNet	5.67x	1.39x	2.82x	39.62x	1.00x

TABLE 4: Memory requirements for the execution of a single iteration of our proposed framework and state-of-the-art techniques. The average memory utilization per dataset is reported (using Python’s memory profiler package), normalized w.r.t the lowest utilization among all methods.

Dataset	Ours	AMC [15]	HAQ [17]	ASQJ [24]	OPQ [18]
CIFAR-10	1.19x	1.08x	1.34x	1.00x	1.15x
CIFAR-100	1.20x	1.09x	1.34x	1.00x	1.15x
ImageNet	1.31x	1.00x	1.50x	1.15x	1.74x

time consuming. Our framework’s overhead is governed by the RL agent’s update scheme, and the reward calculation (mainly its accuracy term), which has to be repeated at each step, for the reasons outlined in Section 4.2.2. The latter leads to elevated time complexity, which is evident for the complex and deep MobileNetV2 architecture. Such calculations are also present in the state-of-the-art approaches, which helps keep the memory utilization measurements at similar levels. So, even though we do not possess an advantage in these comparisons, the overhead introduced by our technique is on par with the state of the art. Note, exploring the joint pruning and quantization design space gives an immediate disadvantage to our proposed technique compared to standalone approaches [15], [17]. Importantly, most related approaches require post-optimization fine-tuning, which significantly increases their memory utilization, especially with the high level of data parallelism normally used for modern datasets. Overhead reduction, as well as improving the scalability of our technique is left for future work.


6 CONCLUSION

In this work, we propose an automated hardware-aware framework to holistically compress DNNs for energy-efficient inference. We explore the joint design space of fine- and coarse-grained pruning and mixed-precision quantization using a novel, composite RL-agent with a custom LUT-based reward. Our framework does not require any fine-tuning and/or long-term retraining steps thus satisfying proprietary and/or privacy constraints. Targeting to reduce energy consumption whilst minimally affecting prediction accuracy, we exceed the state-of-the-art compression techniques (that use fine-tuning). Our framework achieves $\sim 39\%$ average energy reduction across all studied DNNs.

REFERENCES

- [1] N. P. Jouppi *et al.*, “In-datacenter performance analysis of a tensor processing unit,” in *International Symposium on Computer Architecture*, 2017, p. 1–12.

- [2] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 1, pp. 127–138, 2016.
- [3] H. Amrouch, G. Zervakis, S. Salamin, H. Kattan, I. Anagnostopoulos, and J. Henkel, "Npu thermal management," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 11, pp. 3842–3855, 2020.
- [4] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," *Advances in neural information processing systems*, vol. 28, 2015.
- [5] N. Lee, T. Ajanthan, and P. H. Torr, "Snip: Single-shot network pruning based on connection sensitivity," in *International Conference on Learning Representations*, 2019.
- [6] Y. Guo, A. Yao, and Y. Chen, "Dynamic network surgery for efficient dnns," *Advances in neural information processing systems*, vol. 29, 2016.
- [7] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient convnets," in *International Conference on Learning Representations*, 2017.
- [8] X. Gao, Y. Zhao, Ł. Dudziak, R. Mullins, and C.-z. Xu, "Dynamic channel pruning: Feature boosting and suppression," in *International Conference on Learning Representations*, 2019.
- [9] X. Ma *et al.*, "Non-structured dnn weight pruning—is it beneficial in any platform?" *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [10] J. Choi, Z. Wang, S. Venkataramani, P. I.-J. Chuang, V. Srinivasan, and K. Gopalakrishnan, "PACT: Parameterized clipping activation for quantized neural networks," in *International Conference on Learning Representations*, 2018.
- [11] B. Jacob *et al.*, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Conference on Computer Vision and Pattern Recognition*, 2018.
- [12] T.-J. Yang *et al.*, "Netadapt: Platform-aware neural network adaptation for mobile applications," in *European Conference on Computer Vision*, 2018.
- [13] T.-J. Yang, Y.-H. Chen, and V. Sze, "Designing energy-efficient convolutional neural networks using energy-aware pruning," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5687–5695.
- [14] T. Wang *et al.*, "Apq: Joint search for network architecture, pruning and quantization policy," in *Conference on Computer Vision and Pattern Recognition*, 2020.
- [15] Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, and S. Han, "Amc: Automl for model compression and acceleration on mobile devices," in *European Conference on Computer Vision*, 2018.
- [16] G. Zervakis, O. Spantidi, I. Anagnostopoulos, H. Amrouch, and J. Henkel, "Control variate approximation for dnn accelerators," in *Design Automation Conference*, 2021, pp. 481–486.
- [17] K. Wang, Z. Liu, Y. Lin, J. Lin, and S. Han, "Haq: Hardware-aware automated quantization with mixed precision," in *Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8604–8612.
- [18] P. Hu, X. Peng, H. Zhu, M. M. S. Aly, and J. Lin, "Opq: Compressing deep neural networks with one-shot pruning-quantization," in *AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, 2021, pp. 7780–7788.
- [19] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *International Conference on Learning Representations*, 2019.
- [20] W. Niu *et al.*, "Patdnn: Achieving real-time dnn execution on mobile devices with pattern-based weight pruning," in *International Conference on Architectural Support for Programming Languages and Operating Systems*. Association for Computing Machinery, 2020, p. 907–922.
- [21] R. Banner, Y. Nahshan, and D. Soudry, "Post training 4-bit quantization of convolutional networks for rapid-deployment," in *Conference on Neural Information Processing Systems*. Curran Associates Inc., 2019.
- [22] R. Zhao, Y. Hu, J. Dotzel, C. De Sa, and Z. Zhang, "Improving neural network quantization without retraining using outlier channel splitting," *International Conference on Machine Learning*, pp. 7543–7552, June 2019.
- [23] C. Gong, Y. Chen, Y. Lu, T. Li, C. Hao, and D. Chen, "Vecq: Minimal loss dnn model compression with vectorized weight quantization," *IEEE Transactions on Computers*, vol. 70, no. 05, pp. 696–710, 2021.
- [24] H. Yang, S. Gui, Y. Zhu, and J. Liu, "Automatic neural network compression by sparsity-quantization joint learning: A constrained optimization-based approach," in *Conference on Computer Vision and Pattern Recognition*, 2020.
- [25] A. T. Elthakeb, P. Pilligundla, F. Mireshghallah, A. Yazdanbakhsh, and H. Esmailzadeh, "Releq: A reinforcement learning approach for automatic deep quantization of neural networks," *IEEE Micro*, vol. 40, no. 5, pp. 37–45, 2020.
- [26] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," in *International Conference on Learning Representations*, 2015.
- [27] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, "Once-for-all: Train one network and specialize it for efficient deployment," in *International Conference on Learning Representations*, 2019.
- [28] Y. Wang, Y. Lu, and T. Blankevoort, "Differentiable joint pruning and quantization for hardware efficiency," in *European Conference on Computer Vision*. Springer, 2020, pp. 259–277.
- [29] F. Tung and G. Mori, "Deep neural network compression by in-parallel pruning-quantization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 3, pp. 568–579, 2018.
- [30] J. Kim, K. Yoo, and N. Kwak, "Position-based scaled gradient for model quantization and pruning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 20415–20426, 2020.
- [31] S. J. Kwon, D. Lee, B. Kim, P. Kapoor, B. Park, and G.-Y. Wei, "Structured compression by weight encryption for unstructured pruning and quantization," in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1909–1918.
- [32] J. Kim, S. Chang, and N. Kwak, "PQK: Model Compression via Pruning, Quantization, and Knowledge Distillation," in *Inter-speech*, 2021, pp. 4568–4572.
- [33] E. Frantar and D. Alistarh, "Optimal brain compression: A framework for accurate post-training quantization and pruning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 4475–4488, 2022.
- [34] M. Van Baalen *et al.*, "Bayesian bits: Unifying quantization and pruning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5741–5752, 2020.
- [35] Y. He, X. Zhang, and J. Sun, "Channel pruning for accelerating very deep neural networks," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1389–1397.
- [36] H. Pan, X. Niu, R. Li, S. Shen, and Y. Dou, "Dropfilter: a novel regularization method for learning convolutional neural networks," *Neural Processing Letters*, vol. 51, pp. 1285–1298, 2020.
- [37] J.-S. Park *et al.*, "9.5 a 6k-mac feature-map-sparsity-aware neural processing unit in 5nm flagship mobile soc," in *IEEE International Solid-State Circuits Conference*, vol. 64, 2021, pp. 152–154.
- [38] S. Venkataramani *et al.*, "Rapid: Ai accelerator for ultra-low precision training and inference," in *International Symposium on Computer Architecture*, 2021, pp. 153–166.
- [39] N. P. Jouppi *et al.*, "Ten lessons from three generations shaped google's tpuv4i: Industrial product," in *International Symposium on Computer Architecture*, 2021.
- [40] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," in *International Conference on Learning Representations*, 2015.
- [41] M. Hessel *et al.*, "Rainbow: Combining improvements in deep reinforcement learning," in *AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [42] Z. Wang, T. Schaul, M. Hessel, H. van Hasselt, M. Lanctot, and N. de Freitas, "Dueling network architectures for deep reinforcement learning," in *International Conference on Machine Learning*, 2016.
- [43] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "Tetris: Scalable and efficient neural network acceleration with 3d memory," in *International Conference on Architectural Support for Programming Languages and Operating Systems*, 2017, pp. 751–764.
- [44] M. Gao, X. Yang, J. Pu, M. Horowitz, and C. Kozyrakis, "Tangram: Optimized coarse-grained dataflow for scalable nn accelerators," in *International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 807–820.
- [45] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.



Konstantinos Balaskas received his Bachelor Degree in Physics and Master Degree in Electronic Physics from the Aristotle University of Thessaloniki in 2018 and 2020, respectively. Currently, he is pursuing the PhD degree at the same institution, jointly with the Chair for Embedded Systems (CES) at the Karlsruhe Institute of Technology (KIT), Germany. His research interests include machine learning for resource-constrained platforms and physical-driven approximate computing.



Georgios Zervakis is an Assistant Professor at the University of Patras. Before that he was a Research Group Leader at the Chair for Embedded Systems (CES), at the Karlsruhe Institute of Technology (KIT) from 2019 to 2022. He received the Diploma and Ph.D. degrees from the School of Electrical and Computer Engineering (ECE), National Technical University of Athens (NTUA), Greece, in 2012 and 2018, respectively. Dr. Zervakis serves as a reviewer in many IEEE and ACM journals and is also a member of the technical program committee of several major design conferences. He has received one best paper nomination at DATE 2022. His main research interests include low-power design, accelerator microarchitectures, approximate computing, and machine learning.



Andreas Karatzas received the Integrated Master degree (Diploma) from the department of Computer Engineering and Informatics (CEID), University of Patras, Patras, Greece, in 2021. He is currently pursuing the Ph.D. degree at the School of Electrical, Computer and Biomedical Engineering at Southern Illinois University, Carbondale, Illinois, as a member of the Embedded Systems Software Lab. His research interests include embedded systems, approximate computing, and deep learning.



Christos Sad received the Integrated Master degree (Diploma) from the department of Electrical and Computer Engineering (ECE), Aristotle University of Thessaloniki(AUTh), in 2021. He is currently pursuing the Ph.D. degree at the School of Physics, department of Electronics and Computers, Aristotle University of Thessaloniki(AUTh). His research interests include optimization, approximate computing, neural architecture search, GPU programming and GPU acceleration.



Kostas Siozios received his Diploma, Master and Ph.D. Degree in Electrical and Computer Engineering from the Democritus University of Thrace, Greece, in 2001, 2003 and 2009, respectively. Currently, he is Associate Professor at Department of Physics, Aristotle University of Thessaloniki. His research interests include Digital Design, Hardware Accelerators, Resource Allocation and Machine Learning. Dr. Siozios has published more than 170 papers in peer-reviewed journals and conferences. Also, he has

contributed in 5 books of Kluwer and Springer. He has worked as Project Coordinator, Technical Manager or Principal Investigator in 28 research projects funded from the European Commission (EC), European Space Agency (ESA), as well as National Funding.



Iraklis Anagnostopoulos is an Associate Professor at the School of Electrical, Computer and Biomedical Engineering at Southern Illinois University, Carbondale. He is the director of the Embedded Systems Software Lab, which works on run-time resource management of modern and heterogeneous embedded many-core architectures, and he is also affiliated with the Center for Embedded Systems. He received his Ph.D. in the Microprocessors and Digital Systems Laboratory of National Technical University of Athens.

His research interests lie in the area of approximate computing, heterogeneous hardware accelerators, and hardware/software co-design.



Jörg Henkel (M'95-SM'01-F'15) is with Karlsruhe Institute of Technology and was before a research staff member at NEC Laboratories, Princeton, NJ. He has received six best paper awards from, among others, ICCAD, ESWeek and DATE. For two terms he served as the Editor-in-Chief for the ACM Transactions on Embedded Computing Systems. He is currently the Editor-in-Chief of the IEEE Design&Test Magazine and is/has been Associate Editor for major ACM and IEEE Journals. He has led several conferences as a General Chair incl. ICCAD, ESWeek and serves as Steering Committee chair/member for leading conferences and journals for embedded and cyber-physical systems. Prof. Henkel coordinates the DFG program SPP 1500 "Dependable Embedded Systems" and is a site coordinator of the DFG TR89 collaborative research center "Invasive Computing". He is the chairman of the IEEE Computer Society, Germany Chapter, and a Fellow of the IEEE.