

# Decoupled Access-Execute enabled DVFS for tinyML deployments on STM32 microcontrollers

Elisavet Lydia Alvanaki, Manolis Katsaragakis, Dimosthenis Masouros, Sotirios Xydis and Dimitrios Soudris

*Microprocessors and Digital Systems Laboratory, ECE, National Technical University of Athens, Greece*

{ealvanaki, mkatsaragakis, dmasouros, sxydis, dsoudris}@microlab.ntua.gr

**Abstract**—Over the last years the rapid growth Machine Learning (ML) inference applications deployed on the Edge is rapidly increasing. Recent Internet of Things (IoT) devices and microcontrollers (MCUs), become more and more mainstream in everyday activities. In this work we focus on the family of STM32 MCUs. We propose a novel methodology for CNN deployment on the STM32 family, focusing on power optimization through effective clocking exploration and configuration and decoupled access-execute convolution kernel execution. Our approach is enhanced with optimization of the power consumption through Dynamic Voltage and Frequency Scaling (DVFS) under various latency constraints, composing an NP-complete optimization problem. We compare our approach against the state-of-the-art TinyEngine inference engine, as well as TinyEngine coupled with power-saving modes of the STM32 MCUs, indicating that we can achieve up to 25.2% less energy consumption for varying QoS levels.

**Index Terms**—CNN, Microcontroller, Decoupled Access Execution, MCUs, STM32, DVFS

## I. INTRODUCTION

Over the last years, *Edge AI* [1] has appeared to the forefront as a novel computing paradigm. Edge AI refers to the deployment of complex Deep Neural Network (DNN) models directly on edge devices, such as sensors, Internet of Things (IoT) devices, and other embedded systems. Furthermore, the increasing demand for edge-centric digital signal processing applications, such as video analytics [2], mobile visual tasks [3] and others has established Convolutional Neural Networks (CNNs) as the go-to technology for efficiently handling these tasks at the edge. Notably, major technology providers have responded to this paradigm shift by offering enterprise-grade solutions designed to optimize and deploy CNN models at the edge. These solutions include software-based stacks (e.g., Amazon SageMaker Edge [4]) as well as custom, purpose-built hardware chips (e.g., Google's Edge TPU ASIC [5]).

A significant amount of research efforts are focusing on exploring the complex DNN/CNN design space of dataflow schedules [6], [7] to optimize energy. Despite their sophistication, these solutions often assume the presence of hardware accelerators or OS-supported and network-connected edge devices. However, as the computing continuum extends to the far edge of networks, arrays of resource-constrained devices, typified by microcontrollers (MCUs), are introduced into the ecosystem. Unlike their more capable counterparts, MCUs typically employ bare-metal application execution with custom or lightweight run-times and, at times, disconnected from the

network grid. This paradigm shift to the far edge has highlighted the concept of “tinyML” [8]. TinyML involves the optimization of DNN/CNN models to operate seamlessly on these bare-metal MCUs, thus, extending the boundaries of where ML can be applied. In the face of limited processing power and memory, tinyML models are designed to offer efficient inference capabilities, making it possible to bring intelligence to devices previously excluded from the ML landscape.

Deploying CNNs on MCUs introduces two major challenges. First, the limited memory capacity of MCUs poses difficulties in both storing and executing CNN models. This becomes even more complex, since emerging CNN architectures tend to become much deeper, aiming to provide better accuracy and/or support more complex applications [9]. Second, given that MCUs are frequently embedded in battery-operated edge devices, preserving energy resources becomes crucial, since the execution of resource-intensive and computationally hungry DNNs can rapidly deplete the battery, particularly concerning devices with extended operational requirements.

Focusing on CNN inference, prior research [10] and widely-used frameworks (e.g., TFLite Micro [11], Microsoft's NNI [12], ARM's CMSIS-NN [13], etc.) offer optimization techniques that can enable the deployment of tiny CNN models on resource-constrained MCUs. These techniques typically include static, model-specific, post-training optimizations, such as mixed precision arithmetic [14], model's weight pruning [15] and quantization [16], [17], as well as downsampling [18] and Neural Architecture Search (NAS) methodologies [8], [19]. Other works examine schedule and source code optimizations either to enable efficient data and computation mapping on the underlying hardware [20], [21] or to divide CNN layers into independently distributable tasks that can be offloaded and executed in parallel [22].

Although a vast amount of optimizations have been investigated for DNN deployment on MCUs, they mostly focus on models' structural features, e.g., weights, arithmetic, architecture etc., while little attention is given to system-level runtime optimization knobs. In the context of power consumption and/or energy efficiency optimization, Dynamic Voltage and Frequency Scaling (DVFS) [23] forms an efficient tuning knob to balance performance and energy consumption by right-adjusting the clock frequency of the MCU according to the computational demands of the deployed model. Still, the customization of DVFS policy to DNN specific features as well as the materialization of DVFS in practice is not

This work has been partially funded by EU Horizon program CONVOLVE under grant agreement No 101070374 (<https://convolve.eu/>).

straightforward, as it may impose switching overheads [24] not straightforwardly analyzable and increased leakage power due to longer execution times [25].

In this work, we introduce an end-to-end methodology that exploits DVFS for optimizing the energy consumption of CNN inference on low-end MCUs. The proposed methodology is built upon three pillar concepts, i.e. i) selection of energy optimal MCU clocking scheme among differing iso-latency configurations, ii) employment of Decoupled Access Execute (DAE) computing scheme for CNN convolution layers to unlock higher DVFS efficiencies and iii) extraction of DVFS optimal allocation decisions tailored to target CNN architectures and MCU. More specifically, we propose a Decoupled Access Execution scheme, which splits the execution of convolution layers' kernels into memory-bound and compute-bound regions, thus, allowing the exploitation of processor idling during memory accesses. We characterize this new design space for energy-efficient DNNs and we cast/formulate a knapsack optimization problem taking into consideration the MCU's clock parameters and DAE granularities to obtain the optimal DVFS strategies given the neural network architecture and the corresponding QoS constraint. To the best of our knowledge, *this is the first work that examines the application of DVFS for CNN inference on low-end MCUs*. We note that our methodology is orthogonal to optimization frameworks proposed in the past and can be seamlessly applied to pre-/post-training optimized models to further increase the energy efficiency at runtime (e.g., optimized models exported from TFLite Micro [11] or TinyEngine [20]). We showcase the validity and efficacy of our methodology by applying it on top of three CNN inference models derived from the TinyEngine and deploying the resulting model on a STM32 MCU, one of the most popular and widely used MCUs in the embedded systems domain. Our experimental results show that we can achieve up to 25.2% power optimization compared to existing state-of-the-art approaches.

## II. CLOCKING SCHEME OF STM32 MICROCONTROLLERS

The clocking system of STM32 microcontrollers is managed by the Reset and Clock Control (RCC) peripheral. The RCC provides a wide range of clocks and clock sources which cater to various system requirements, e.g., peripheral Bus and UART clocks [26]. In this work, we focus on specific clocks and settings, which determine the frequency of the system clock (SYSCLK) of the MCU, responsible for driving the CPU core, memory, and some peripheral modules. Figure 1 illustrates the simplified circuit diagram of how the SYSCLK is determined. Specifically, the output frequency of the clock can be configured by the following clock sources and settings:

- **High-Speed Internal (HSI) Clock:** The HSI clock source is an internal oscillator within the MCU. The HSI clock operates at 16MHz by default. SYSCLK can be derived directly from HSI or through the PLL when HSI is selected as the PLL input source.
- **High-Speed External (HSE) Clock:** The HSE clock source is an external clock provided by an external crystal oscillator or clock generator. Depending on the MCU, the HSE clock can be configured to run at various frequencies, with our examined board supporting a range from 1 to 50MHz. Similar

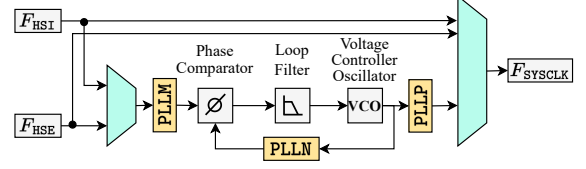


Fig. 1: Simplified circuit diagram for clock configuration through HSE and PLL parameters.

to the HSI, the SYSCLK can be derived directly from HSE or through the PLL when HSE is selected as the PLL input.

- **Phase-Locked Loop (PLL):** The PLL is a hardware module which allows to multiply the frequency of the selected input clock source (either HSI or HSE) by a programmable factor to generate a higher-frequency output. As shown in Fig. 1, this programmable factor is determined by input and output dividers within the PLL circuit, which simplify the PLL design and achieve stability requirements within a wider range of input and output frequencies [26]. Specifically, the frequency of the system clock can be calculated by the following equation:

$$F_{\text{SYSCLK}} = F_{\{\text{HSE}, \text{HSI}\}} * \frac{\text{PLLN}}{\text{PLLM} * \text{PLLP}} \quad (1)$$

where, PLLM is a factor that determines how much the input frequency is multiplied before it reaches the Voltage-Controlled Oscillator (VCO) in the PLL; PLLN is the multiplication factor for the VCO input frequency to determine the VCO output frequency, which is then provided as feedback to the Phase Comparator, aiming to provide input/output synchronization. The loop filter is utilized in order to ensure system stability and mitigate ripple effects during clock startup; Last, PLLP determines the division factor to obtain the final SYSCLK frequency.

### A. Considerations of SYSCLK Frequency Scaling

Determining the frequency of the system's clock,  $F_{\text{SYSCLK}}$ , can be achieved in different ways (e.g., directly through the HSI/HSE clock sources or through the PLL circuit). Choosing between these alternatives involves important trade-offs. To examine the impact of each alternative on the operational efficiency and power consumption of the MCU, we develop and execute a specialized microbenchmark, designed to execute repetitive addition operations within a loop. We focus our exploration specifically on the HSE and PLL parameters, since the HSI clock source yields higher power consumption compared to the HSE and is also prone to drift and jitter, thus, providing less stability and precision [26]. We set the value of the PLLP equal to 2, which is the minimum possible value for the divider, since for the same  $F_{\text{SYSCLK}}$ , selecting a higher PLLP value leads to a higher required VCO frequency and, thus, higher power consumption (as evident from Fig. 1 and Eq. 1).

**Power consumption of iso-frequency configurations:** Figure 2 illustrates the impact of different HSE, PLLM and PLLN configurations on the power consumption of the board for different SYSCLK frequencies. Two major observations are derived from Fig. 2: (i) the same output frequency can be generated through different HSE, PLLM

and PLLN combinations, however (ii) the selected configuration strongly affects the power consumption on the STM32 MCU. For instance, generating an SYSCLK output of 100MHz for the tuples  $\{50, 25, 216\}$  and  $\{16, 8, 100\}$  leads to 50% power gap. The combinations that minimize the power consumption are selected for the target SYSCLK. Different combinations may generate the same output frequency and power consumption, e.g.  $\{50, 25, 100\}$  and  $\{50, 50, 200\}$ , thus further investigation is required for selecting the optimal combination. Similar observations are derived for other SYSCLK configurations.

#### Switching between different SYSCLK frequencies:

Generating the SYSCLK frequency using the PLL module introduces a notable switching overhead ( $\approx 200\mu\text{sec}$ ), since, when modifying the PLL parameters, the circuit has to be restarted, resulting in a substantial delay per switch. On the other hand, switching from the PLL frequency to the HSE clock occurs almost instantly, due to the direct wiring of the HSE with the SYSCLK (Fig. 1). Thus, for high-to-low (i.e.,  $< 50\text{MHz}$ ) frequency switches, opting for the HSE clock over re-calibrating the PLL parameters can be beneficial.

### III. DECOUPLED ACCESS EXECUTE ENABLED DVFS ON STM32 MCUS

This section introduces the proposed methodology for optimizing energy consumption of CNN model inference on a STM32 MCU under latency constraints (a.k.a. QoS). The rationale behind our approach is that end-users often have distinct inference latency/throughput constraints for their applications and/or even operating in the context of battery-operated far-edge MCUs. Figure 3 shows an end-to-end overview of the proposed approach. The proposed methodology consists of three distinct phases (described in Sections III-A - III-C), applicable to both unoptimized CNN models and optimized ones, e.g., models exported from TinyEngine [20].

#### A. Step 1: Memory Access & CPU Execution Decoupling

The first step of our proposed methodology focuses on source-code level restructuring, with the objective to Decouple memory Accesses from CPU Execution (DAE), thus, creating memory-bound and compute-bound sub-segments within the layer's structure. DAE restructuring is a key-enabler for our strategy, as it provides more fine-grained control over when and how frequently memory accesses and computations occur (Sec. III-B). This enables the application of different frequencies for extended durations and with finer control, tailored to the specific requirements of each operation, such as memory access and computation, thus, resulting in the mitigation of clock switching overhead issues. We first identify the CNN model's most computationally-intensive and time-consuming

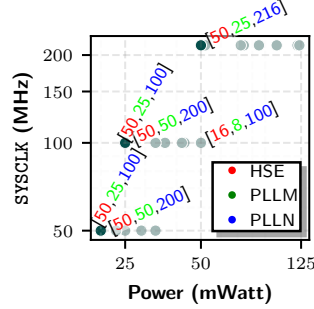


Fig. 2: Clock Frequency and Power for different HSE, PLLM and PLLN configurations.

```

1 for (channel = 0; channel < in_channels/g; channel+=g) {
2     //LFO for memory bound operations (Sec.IIIB)
3     ClockSwitchHSE(hse);
4     //Memory Bound Segment: Load g channels from the
5     //feature maps
6     q15_t buf1, ..., bufg = getChannels(ch1, ..., chg);
7     //HFO for computation (Sec. IIIB)
8     ClockSwitchPLL(pllm, plln, hse);
9     // Compute Bound Segment: Perform depthwise
10    //convolution for each channel
11    convolve_depthwise(kernel, buf1, **args);
12    convolve_depthwise(kernel, buf2, **args);
13    ...
14    convolve_depthwise(kernel, bufg, **args);
15 }

```

Listing 1: Simplified source code modification for enabling decoupled access-execution (DAE) in depthwise convolutions.

layers (1A). We focus and apply DAE (1B) on two specific layer types, i.e., *i*) depthwise and *ii*) pointwise convolutions. These layer types make up over 80% of the total number of layers found in deep lightweight CNN models, e.g., Mobilenet [3], which employ the concept of depthwise separable convolution to reduce model's size and complexity.

**Depthwise Convolutions:** Depthwise convolution is a specialized CNN operation where each input channel is convolved with a separate learnable filter, capturing spatial features per channel. Typically, CNNs gradually learn increasingly complex features, each one represented by a different channel. For instance, in an image, the initial 3 input channels (RGB) increase as the network processes the image to extract and represent more complex features, such as textures, specific objects and others. State-of-the-art frameworks, such as CMSIS-NN [13] and TinyEngine [20] implement a per-channel computation approach for depthwise convolutions. In contrast, our DAE approach introduces a parametric unrolling factor called the "decoupling granularity", denoted by  $g$ . This factor determines the number of channels that are buffered in cache memory before the convolution operation is executed on each of them, thus, separating the code into memory-bound and compute-bound subsegments. Listing 1 provides a simplified code snippet that illustrates the practical implementation of DAE optimization, showing how the decoupling granularity enables the efficient execution of depthwise convolutions. For example, when  $g = 4$ , four channels are fetched in the MCU's cache memory before proceeding to the actual computation. This division enables the application of different clock frequencies per segment, which we further discuss in Sec. III-B.

**Pointwise Convolutions:** Depthwise convolution is often followed by pointwise convolution to perform channel-wise mixing and dimensionality reduction, thus, reducing model size and computational complexity while maintaining performance. Pointwise convolutions involve  $1 \times 1$  kernel sizes and are applied to each element within input channels. CMSIS-NN[13] and TinyEngine [20] implement pointwise convolutions in a per-column manner. Each column consists of one element per input channel. Our approach performs decoupling on the per channel memory accesses, thus splitting the code segment into memory and compute-bound regions. Just as in the depthwise convolution, we introduce the concept of the decoupling granularity, denoted as  $g$ , for modular buffering support w.r.t. the

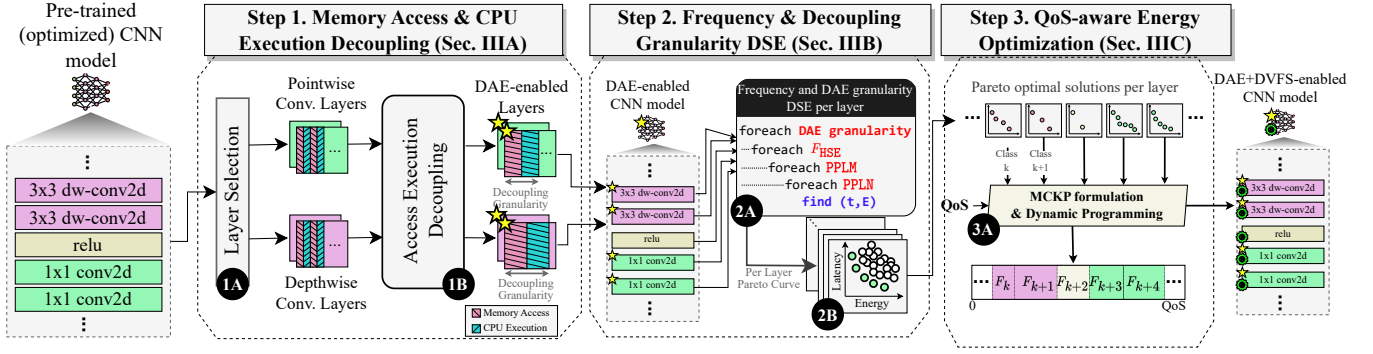


Fig. 3: Overview of the proposed methodology.

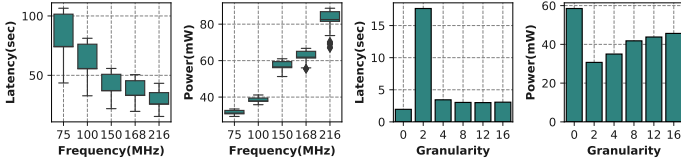


Fig. 4: Impact of different DAE and clocking configurations on latency and power of depthwise and pointwise layers.

number of columns fetched from memory. For instance, for a  $8 \times 8 \times 3$  input image and a  $1 \times 1 \times 3$  kernel,  $g$  columns are loaded into the cache before performing computation for each one, in contrast with TinyEngine and CMSIS-NN, which load a single  $1 \times 1 \times 3$  image column at a time. In general, integrating multiple buffers leads to the generation of larger memory/compute-bound regions, thus we can minimize the frequency switching overhead while avoiding high power consumption. However, very high buffer size can lead the cache misses to skyrocket resulting in performance degradation.

After the DAE phase is performed, the CNN layers are encapsulated with a configurable decoupling granularity factor for the layer unrolling. The modified DAE-enabled CNN model is propagated to the next step of our proposed methodology (Step 2.) for effective exploration and configuration of the DAE granularity factor, alongside with the DVFS parameters exploration. DAE-enabled CNNs entail no accuracy drops.

#### B. Step 2: DAE and Clocking Co-exploration

At this step, we analyze the performance and energy consumption in a per-layer manner (2A) considering the interplay effects of DAE and clocking scheme configurations. To measure the energy consumption and performance of each layer, we have developed and integrated a custom run-time monitoring mechanism for supporting per-layer monitoring and profiling. Our mechanism relies on the on-board timers of the target MCU, which are triggered in-between the layers' code segments. Furthermore, we take advantage of STM32 MCUs integrated support for power sampling and we monitor the power consumption prior and after the DVFS integration on every CNN layer. The power and performance metrics for the DVFS in each layer are aggregated and utilized for the design space exploration for DAE and clocking configuration.

More specifically, we co-explore and gain insights on the design space defined by the following three key parameters: *i*)

the decoupling granularity factor  $g$  described in Sec. III-A; *ii*) the clock frequency of the SYSCLK clock; and *iii*) the selection of parameters for the PLL module (namely PLLM and PLLN), both described in Sec. II. Regarding the decoupling granularity, determining the most suitable value per layer depends on both board-related specifications (e.g., cache size) as well as code-related characteristics (e.g., number of output channels and kernel size). In our case, we examine six different values, i.e.,  $g \in \{0, 2, 4, 8, 12, 16\}$ , where  $g = 0$  indicates no DAE optimization and corresponds to the default input model.

Regarding the different clock alternatives, we define two parametric operating modes, namely the *Low Frequency Operation* (LFO) and the *High Frequency Operation* (HFO). LFO exclusively employs the HSE clock source at a predefined frequency (50MHz) and aims to reduce power consumption on the board. In contrast, HFO configures the system's clock using the PLL circuit, where the final SYSCLK value is determined by varying combinations of the PLLN and PLLM parameters, specifically  $PLLN \in \{75, 100, 150, 168, 216, 336, 432\}$  and  $PLLM \in \{25, 50\}$ . This distinction between the two operating modes allows us to quickly transition between different SYSCLK frequencies, thus, mitigates susceptibility to the switching overhead associated with the PLL, as elaborated in Sec. II-A. Overall, DVFS switching is performed between the memory (Lst. 1:3) and compute (Lst. 1:7) bound regions, in order to exploit the decoupled access-execution transformation optimally, with LFO applied to the memory-bound and HFO to the compute-bound subsegments respectively. A similar approach is followed in pointwise convolution layers.

**DSE Insights:** Figure 4 shows the impact of varying operating frequencies (left) and granularity factors  $g$  (right) on the layer latency and power consumption. First, we observe that as the number of operating frequency increases, the power consumption is traded for better performance, thus composing the design space. Moreover, changing values of the granularity factor can provide significant variation on the latency and the power consumption. For instance, power consumption can drop up to 54.2% compared to the initial execution. Thus, the effective co-exploration of granularity  $g$  and frequency leads to power/latency trade-offs. The result of the DSE is a solution space per layer, where each solution trade-offs between performance and power consumption (2B). In this space, we



select the Pareto optimal points, to be propagated to Step 3.

### C. Step 3: QoS-aware Energy Optimization

In the final stage, we determine the optimal frequencies for each layer within the CNN, aiming to minimize the model's total energy consumption while satisfying a predefined latency budget (QoS). We denote the set of all possible frequencies generated either through the PLL or the HSE clock as  $F$  and the set of all possible granularity factors as  $G$ . Let  $n$  be the total number of layers of the CNN model and  $P_k = \{\dots, p_j^k = \{t_j^k, E_j^k\}, \dots\}$ ,  $k \in \{1, \dots, n\}$ ,  $j \in \{1, \dots, |P_k|\}$  be the set of Pareto optimal solutions (from Step 2) of layer  $k$ , where  $t_j^k$  and  $E_j^k$  denote the latency and the energy consumption of the  $j^{th}$  pareto optimal solution for layer  $k$  when operating with DVFS enabled with an HFO frequency  $f \in F$  and granularity  $g \in G$ . We consider the minimization of the overall energy  $E$  of the CNN deployed on the target STM32 MCU, so that the overall execution time  $T$  does not go beyond a user-defined QoS. Then, the target optimization problem can be formulated as follows:

$$\text{minimize } E = \sum_{k=1}^n \sum_{j \in P_k} E_j^k x_{kj} \quad (2)$$

$$\text{s.t. } T = \sum_{k=1}^n \sum_{j \in P_k} t_j^k x_{kj} \leq QoS \quad (3)$$

$$\sum_{j \in P_k} x_{kj} = 1, \quad k = 1, \dots, n \quad (4)$$

$$x_{kj} \in \{0, 1\}, \quad k = 1, \dots, n, j \in P_k. \quad (5)$$

We model our problem according to the Multiple-Choice Knapsack Problem (MCKP) [27], which extends the classical knapsack problem by categorizing items into distinct classes (3A). In this formulation, the binary decision of including an item is replaced by the selection of precisely one item from each class and the goal is to maximize the value of items included in the knapsack while not exceeding its size. In our case, each individual class represents the various Pareto optimal solutions  $p_j^k$  per layer  $k$ . Each item in the class is characterized by its own value (i.e., energy consumption  $E_j^k$ ) and size (i.e., latency  $t_j^k$ ). Our goal is to minimize the overall energy consumption ( $E$ ) while adhering to the size constraint ( $T < QoS$ ). We convert our minimization objective to a maximization one using the transformation found in [27]. Last, we solve the optimization problem using a pseudo-polynomial time solution based on a dynamic programming (DP) approach.

## IV. EXPERIMENTAL SETUP AND EVALUATION

**Experimental Setup:** Our experimental evaluation is conducted on an STM32F767ZI Nucleo board, equipped with an ARM Cortex M7 CPU featuring a 16KB L1-cache. The board incorporates a High-Speed External (HSE) clock, ranging from 1MHz to 50MHz. For power consumption monitoring, we employed the INA219 power sensor. To mitigate potential variations arising from temperature-induced power fluctuations, we systematically compared each power measurement with the power consumption of the baseline input model at the corresponding timestamp. Our proposed methodology is evaluated over three pre-trained CNN models, namely Visual

Wake Words (VWW), Person Detection (PD), and Mobilenet-V2 (MBV2), derived from the MCUNet [20] inference library. Hardware-Aware Neural Architecture Search and linear int8 quantization have been used for quantization. We conduct a comparative analysis between our approach and the state-of-the-art TinyEngine [20], which serves as the baseline for evaluation. Our experiments are conducted in an iso-latency execution scenario, where we measure energy consumption for a certain period specified by a QoS constraint. In the case of TinyEngine, this entails the board remaining in an idle state with a constant frequency of 216MHz after an inference, until the QoS threshold is met. We also consider TinyEngine enhanced with clock gating, a technique designed to optimize power consumption by selectively deactivating non-utilized board clocks and the voltage regulator, thus minimizing power leakage throughout the CNN inference.

**Energy Comparison and Analysis:** Figure 5 illustrates a comparison of energy consumption between our proposed approach, and the two configurations of the TinyEngine: one without any optimization and the other with clock gating applied. This evaluation encompasses various CNN inference models, each subject to discrete Quality of Service (QoS) constraints set at 10% (tight), 30% (moderate), and 50% (relaxed). X-axis illustrates the QoS constraints, while Y-axis represents the normalized energy consumption. Our proposed approach surpasses both instances of the TinyEngine, exhibiting a reduction in energy consumption up to 25.2%. Additionally, compared to the TinyEngine equipped with clock gating we achieve up to 7.2% less energy consumption. Furthermore, our observations indicate that relaxing the QoS constraints can lead to a notable reduction in energy consumption, albeit at the cost of some performance trade-offs. For instance, when examining the Mobilenet-V2 model, the energy consumption of our approach under a relaxed 50% QoS constraint decreases to 20.4% compared to the stringent 10% constraint.

**Frequency Scaling Analysis:** Figure 6 illustrates the HFO for each examined CNN. X axis indicates the corresponding layer type as the CNN execution proceeds and the granularities selected for 10% and 50% QoS constraint, respectively, while Y axis shows the operating frequency per layer. The LFO configuration at 50MHz of the memory-bound segment of each layer is excluded for simplicity. The observations derived are the following. Firstly, the operational frequency is configured to maximum (216MHz) mostly for performing

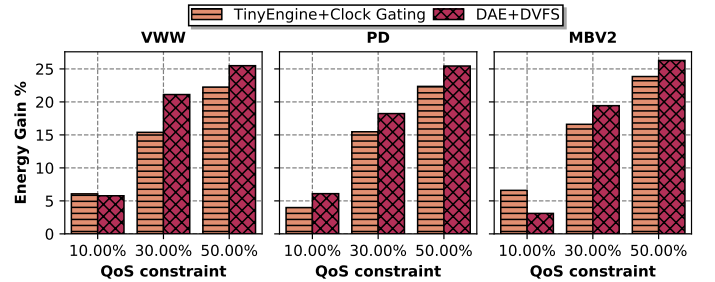


Fig. 5: Energy consumption gains of our approach over the TinyEngine [20] baseline. We compare against TinyEngine with Clock Gating over the examined CNN models.

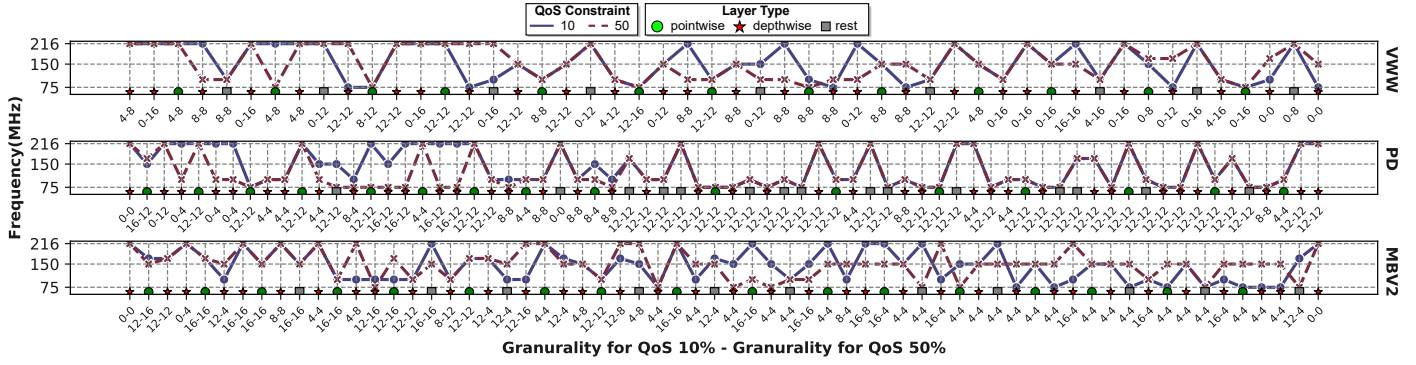


Fig. 6: Frequency distribution throughout layer progression over the examined CNN models for 10% and 50% QoS constraints.

pointwise convolutions, i.e. 58.8% against 21.4% for depthwise convolutions. The latter are less compute-intensive, thus decreasing the operational frequency will not lead to significant performance degradation. Furthermore, the 46.1% of the pointwise convolutions and 43.4% of the depthwise convolutions are executed over the lowest operating frequencies, i.e. 75MHz and 100MHz, aiming to boost the optimization objective of power minimization. Last but not least, we investigate the impact of QoS constraints on the operating frequency. Our experiments indicate that 18.6% more layers are operating at 216MHz for tight constraints (10%). Regarding the granularity analysis, for the relaxed QoS(50%), 22.3% more layers operate with granularity factor 16, compared to 10% constraint. This is due to the fact that there is higher space to trade latency, thus the computation-bound parts are split to bigger segments, aiming to minimize switching overhead and provide power reduction.

## V. CONCLUSION

In this work, we present a novel end-to-end methodology that exploits DVFS for optimizing the energy consumption of CNN inference on low-end STM32 MCUs. Our approach strongly leverages Decoupled Access Execute techniques to discretize memory-bound and compute-bound layer segments. Our approach outperforms existing state-of-the-art approaches by achieving up to 25.2% less energy consumption.

## REFERENCES

- [1] E. Li, Z. Zhou, and X. Chen, "Edge intelligence: On-demand deep learning model co-inference with device-edge synergy," in *Proceedings of the 2018 Workshop on Mobile Edge Communications*, pp. 31–36, 2018.
- [2] G. Ananthanarayanan, P. Bahl, P. Bodík, K. Chintalapudi, M. Philipose, L. Ravindranath, and S. Sinha, "Real-time video analytics: The killer app for edge computing," *computer*, vol. 50, no. 10, pp. 58–67, 2017.
- [3] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint*, 2017.
- [4] "Amazon sagemaker edge." <https://aws.amazon.com/sagemaker/edge/>.
- [5] "Google edge tpu." <https://cloud.google.com/edge-tpu>.
- [6] L. Mei, P. Houshmand, V. Jain, S. Giraldo, and M. Verhelst, "Zigzag: Enlarging joint architecture-mapping design space exploration for dnn accelerators," *IEEE Transactions on Computers*, 2021.
- [7] H. Kwon, P. Chatarasi, V. Sarkar, T. Krishna, M. Pellauer, and A. Parashar, "Maestro: A data-centric approach to understand reuse, performance, and hardware cost of dnn mappings," *IEEE micro*, 2020.
- [8] C. Banbury, C. Zhou, I. Fedorov, R. Matas, U. Thakker, D. Gope, V. Janapa Reddi, M. Mattina, and P. Whatmough, "Micronets: Neural network architectures for deploying tinyml applications on commodity microcontrollers," *Proceedings of Machine Learning and Systems*, 2021.
- [9] A. K. Kakolyris, M. Katsaragakis, D. Masouros, and D. Soudris, "Road-runner: Collaborative dnn partitioning and offloading on heterogeneous edge systems," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1–6, IEEE, 2023.
- [10] K. T. Chitty-Venkata and A. K. Somani, "Neural architecture search survey: A hardware perspective," *ACM Computing Surveys*, 2022.
- [11] R. David, J. Duke, A. Jain, V. Janapa Reddi, N. Jeffries, J. Li, N. Kreeger, I. Nappier, M. Natraj, T. Wang, et al., "Tensorflow lite micro: Embedded machine learning for tinyml systems," *Proceedings of Machine Learning and Systems*, vol. 3, pp. 800–811, 2021.
- [12] "Microsoft nni." <https://github.com/microsoft/nni>.
- [13] L. Lai, N. Suda, and V. Chandra, "Cmsis-nn: Efficient neural network kernels for arm cortex-m cpus," *arXiv preprint arXiv:1801.06601*, 2018.
- [14] A. Capotondi, M. Rusci, M. Fariselli, and L. Benini, "Cmix-nn: Mixed low-precision cnn library for memory-constrained edge devices," *IEEE Transactions on Circuits and Systems II: Express Briefs*, 2020.
- [15] J. D. De Leon and R. Atienza, "Depth pruning with auxiliary networks for tinyml," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2022.
- [16] J. Moosmann, H. Mueller, N. Zimmerman, G. Rutishauser, L. Benini, and M. Magno, "Flexible and fully quantized ultra-lightweight tinyssimoyolo for ultra-low-power edge systems," *arXiv preprint:2307.05999*, 2023.
- [17] O. Spantidi and I. Anagnostopoulos, "Automated energy-efficient dnn compression under fine-grain accuracy constraints," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023.
- [18] O. Saha, A. Kusupati, H. V. Simhadri, M. Varma, and P. Jain, "Rnnpool: Efficient non-linear pooling for ram constrained inference," *Advances in Neural Information Processing Systems*, vol. 33, pp. 20473–20484, 2020.
- [19] E. Liberis, L. Dudziak, and N. D. Lane, "μnas: Constrained neural architecture search for microcontrollers," in *Proceedings of the 1st Workshop on Machine Learning and Systems*, pp. 70–79, 2021.
- [20] J. Lin, W.-M. Chen, Y. Lin, C. Gan, S. Han, et al., "Mcnunet: Tiny deep learning on iot devices," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11711–11722, 2020.
- [21] S. Jaiswal, R. K. K. Goli, A. Kumar, V. Seshadri, and R. Sharma, "Minun: Accurate ml inference on microcontrollers," in *Proceedings of the 24th ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems*, pp. 26–39, 2023.
- [22] Z. Zhao, K. M. Barijough, and A. Gerstlauer, "Deepthings: Distributed adaptive deep learning inference on resource-constrained iot edge clusters," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 11, pp. 2348–2359, 2018.
- [23] D. C. Snowdon12, E. Le Sueur, S. M. Petters, and G. Heiser, "A platform for os-level power management," *The European Professional Society on Computer Systems 2009*, 2009.
- [24] B. Acun, K. Chandrasekar, and L. V. Kale, "Fine-grained energy efficiency using per-core dvfs with an adaptive runtime system," in *2019 Tenth International Green and Sustainable Computing Conference (IGSC)*, pp. 1–8, IEEE, 2019.
- [25] R. Jejurikar, C. Pereira, and R. Gupta, "Leakage aware dynamic voltage scaling for real-time embedded systems," in *Proceedings of the 41st annual Design Automation Conference*, pp. 275–280, 2004.
- [26] W. Gay, "Beginning stm32," *Beginning STM32*, 2018.
- [27] H. Kellerer, U. Pfersch, D. Pisinger, H. Kellerer, U. Pfersch, and D. Pisinger, "The multiple-choice knapsack problem," *Knapsack Problems*, pp. 317–347, 2004.