

DynaMIX: Resource Optimization for DNN-Based Real-Time Applications on a Multi-Tasking Platform

Minkyung Cho and Kang G. Shin
The University of Michigan, Ann Arbor, MI, U.S.A.
{minkycho, kgshin}@umich.edu

Abstract—As deep neural networks (DNNs) prove their importance and feasibility, more and more DNN-based apps, such as detection and classification of objects, have been developed and deployed on autonomous vehicles (AVs). To meet their growing expectations and requirements, AVs should “optimize” use of their limited onboard computing resources for multiple concurrent in-vehicle apps while satisfying their timing requirements (especially for safety). That is, real-time AV apps should share the limited on-board resources with other concurrent apps without missing their deadlines dictated by the frame rate of a camera that generates and provides input images to the apps. However, most, if not all, of existing DNN solutions focus on enhancing the concurrency of their specific hardware without dynamically optimizing/modifying the DNN apps’ resource requirements, subject to the number of running apps, owing to their high computational cost. To mitigate this limitation, we propose DynaMIX (Dynamic MIXed-precision model construction), which optimizes the resource requirement of concurrent apps and aims to maximize execution accuracy. To realize a real-time resource optimization, we formulate an optimization problem using app performance profiles to consider both the accuracy and worst-case latency of each app. We also propose dynamic model reconfiguration by lazy loading only the selected layers at runtime to reduce the overhead of loading the entire model. DynaMIX is evaluated in terms of constraint satisfaction and inference accuracy for a multi-tasking system and compared against state-of-the-art solutions, demonstrating its effectiveness and feasibility under various environmental/operating conditions.

I. INTRODUCTION

As more deep neural network (DNN) apps are getting added in autonomous vehicles (AVs), it is critically important to ensure the timeliness and quality of their execution. Major car-makers, like Tesla [1] and Toyota [2], are expanding the scope of using DNNs for camera-based vision apps. In AVs, a camera continuously monitors environmental conditions at a rate of 10 – 40 frames per second (FPS) [3], [4]. Since the results of processing the current image frame should be produced before the next frame arrives, all in-vehicle apps using the image frames captured by a camera must meet the same deadline dictated by the camera’s frame rate. However, the interference between concurrent vision apps may lead to miss some of their deadlines. Fig. 1 shows a motivating scenario in which two apps process image frames from the same camera. In such a case, the two apps have the *same* deadline since they must neither produce outdated data nor interfere with the processing of the next image frames. If the execution of a classification




Fig. 1. Example use-case and the objective of DynaMIX

app precedes that of a real-time object detection (RTOD) app, the end-to-end (i.e., camera-to-detection output) latency of the object detection app exceeds its deadline.

However, guaranteeing the deadline of embedded control apps like AVs remains a challenging problem. When the number of concurrent apps changes dynamically, the amount of resources (e.g., CPU and memory space) they require should also change to meet their deadline. Unfortunately, most vision apps are computation-intensive, and hence state-of-the-art (SOTA) multi-tasking systems [3], [12]–[14] focused on concurrent execution of multiple DNN apps on specific processors without dynamically modifying their resource requirements. The most straightforward way to reduce the amount of resources required by each app is compressing the DNN model. Given the tradeoff between accuracy and computation, uniformly compressing all the layers in a model with a single precision is not efficient. For example, quantizing a model with high precision may result in a deadline miss. Conversely, low precision quantization will deteriorate accuracy [15], which may cause fatal accidents. This tradeoff calls for a design that optimizes each DNN model to be represented in mixed-precision while considering *total accuracy* (defined as the weighted sum of accuracies of concurrent apps) as well as its deadline.

In this paper, we propose DynaMIX (Dynamic MIXed-precision model construction), which dynamically changes the bit-width setting of each DNN model according to the status of running apps to optimize their computation and memory resources. For the realization of DynaMIX, we must address two difficulties that make mixed-precision models unsuitable for time-critical systems like AVs. First, we have to determine the optimal degree of compression (i.e., optimal bit configuration) for each DNN model at runtime. Balancing

TABLE I
COMPARISON WITH PRIOR WORK. “MULTI-TASKING SUPPORT” DENOTES RUNNING VARIOUS APPS NOT CONSTRAINED BY SPECIFIC DNN MODELS

	Multi-Tasking Support	Real-Time Guarantee	No Binding to Specific Processor
A. Resource Management [5]–[7]	✓		✓
B. RTOD Execution [3], [8]–[10]		✓	✓
C. Hardware-based Scheduling [11]–[14]	✓	✓	
DynaMIX	✓	✓	✓

between latency and accuracy at runtime is difficult as both require to run DNN inferences as many as the number of compressed models of each app when the status of running apps changes, thus incurring high online overhead. To address this difficulty, we create measurements-based app performance profiles via regression for worst-case latency and most-likely inference accuracy according to the available memory size. The resulting profiles are then used for online resource optimization across multiple apps.

Second, we need to reduce the overhead of model loading. When their optimal bit settings are modified, loading the entire mixed-precision models is inefficient unless the bit-widths of all their layers are changed. However, loading *all* possible mixed-precision models generated from one DNN model into memory is neither efficient nor possible. To reduce this model switching overhead, we propose runtime model reconfiguration based on the concept of lazy loading, commonly used in other domains such as web services and operating systems [16], [17]. The proposed layer-wise lazy loading helps avoid the need for loading the entire model whenever the optimal bit setting is changed to use the platform resources efficiently.

DynaMIX is composed of offline and online stages. In the online stage, it formulates the resource optimization as a constrained nonlinear optimization problem. The constraints of the optimization include not only the deadline but also the offline-generated profiles to determine an optimal solution (i.e., memory space for each app) that maximizes the total accuracy subject to the deadline and the available memory space. When a new camera image enters the system and any change in the status of concurrent apps is detected, the main process adjusts the amount of memory space allocated to the concurrent apps. After determining the new bit configuration with the adjusted memory space for each DNN app, the app reconfigures its execution path by loading the necessary high-precision layers to optimize the amount of computation resource. This way, all of them can meet the deadline without compromising the total accuracy.

We have evaluated DynaMIX for running multiple concurrent apps. Our findings help one or more concurrent apps meet the deadline while optimizing the resource allocation. The resulting model for each app is shown to lose $< 2.2\%$ accuracy even in the case of an increased number of concurrent apps or tight timing constraints. Furthermore, we have shown the feasibility of DynaMIX’s resource allocation in various aspects and demonstrated its effectiveness in handling multi-

tasking scenarios. We have also identified the cases DynaMIX could not support, or the resulting total accuracy was too low to be acceptable for the apps in reality.

In summary, this paper makes the following contributions:

- The first resource optimization scheme for concurrent in-vehicle apps to meet their deadline without binding to specific processors;
- Formulation of the optimization problem using app performance profiles to find the optimal compression degrees for multiple apps while accounting for their accuracies and worst-case latencies in real time;
- Runtime model reconfiguration while lazy loading only the selected layers when the optimal bit setting of an app varies, to mitigate the entire model loading overhead;
- Facilitation of rapid deployment of existing DNN apps without any (re)training required to prevent any significant accuracy drop in runtime adaptive quantization.

II. RELATED WORK

Prior work on multiple or real-time DNN inferences on embedded platforms has evolved into three main branches: resource management, RTOD optimization, and hardware-based scheduling (Table I).

A. Resource Management for Multiple Vision Apps

Resource management plays a critical role in ensuring the continuous execution of multiple vision apps. Most resource-management schemes reduce the excessive cost of DNN pipeline execution by fully/partially offloading the computation to a cloud server or other edge devices. MCDNN [5], a representative resource-management scheme, proposes a runtime scheduler to process DNN models (or fragments thereof) across both the edge device(s) and the cloud server. Although offloading reduces the computation load of the edge system, it yields inconsistent app quality because of unpredictable cloud accessibility or possible privacy leakage.

To mitigate/avoid these problems, [6], [7] suggested on-device resource management. DeepMon [6] aims to guarantee continuous vision apps by optimizing the convolutional neural networks (CNN) on mobile GPUs. It accelerated the convolution by reusing the intermediate results via caching. NestDNN [7] proposed a filter pruning for resource management. However, both DeepMon and NestDNN support only non-real-time tasks, probably because the most commonly used real-time tasks (e.g., object detection or tracking) require significant

amounts of computation that the edge system cannot complete in a timely manner (e.g., DeepMon shows only 1~2 FPS). DynaMIX was inspired by [6], [7] in that both pursue a serverless method. DynaMIX reduces the complexity of vision apps via DNN compression, and finds the optimal models by exploiting the knowledge of app performance.

1) *Why not adapting NestDNN [7] to guarantee deadlines?*: Even if NestDNN’s scheduler is revised to meet real-time constraints, it cannot guarantee both accuracy and timeliness. NestDNN leverages a small, fixed number (i.e., five) of pruned models for each app, thus sacrificing accuracy to create models of various sizes. If more pruned models were generated for accuracy, timely execution is difficult because of its high overhead of finding an optimal model set for concurrent DNN apps. To address this, they need to formulate an optimization problem using app performance profiles and generate lookup tables as we proposed here. In contrast, DynaMIX fully supports real-time scenarios.

B. RTOD Execution on Embedded Systems

The great potential of RTOD for improved safety and convenience has yielded a large body of work running RTOD on embedded systems, particularly AVs. Despite the significant recent advances in object detection [18]–[20], RTOD still remains to be a bottleneck in embedded real-time systems due to its heavy computation requirement to localize and classify multiple objects captured in a stream of camera images.

AdaVP [8] built a parallel object detection and tracking pipeline to run two apps in parallel; when a new object is detected, the DNN setting (i.e., input frame size) is adjusted at runtime to increase the tracking accuracy. Yang *et al.* [3] handled simultaneous vision apps with private camera streams and identical DNN models; throughput was increased by sharing the base architecture and processing multiple camera images with multiple threads. DNN-SAM [10] enabled multiple inferences with one RTOD model to process the images produced by different cameras. It ensures the deadline by splitting the original image into different portions according to the criticality levels and adjusts the size of each portion via scaling.

The authors of [9] proposed a new RTOD system that could change the execution path based on a *dynamic deadline* — the deadline of a RTOD task varies with the underlying driving environment, like skipping layers or choosing one of the sub-networks of different sizes. Moreover, they determined the execution path of RTOD using per-layer latency (for all layers) and dynamic deadline, but did not account for the interruptions by other concurrent apps in their path-selection decisions. DynaMIX was inspired by [9] in that the execution path can be changed dynamically to account for the execution environmental condition. DynaMIX reconfigures the DNN execution pipeline once the amount of available resources for each app changes.

In summary, prior work accelerates the RTOD task to meet its explicit/implicit deadline using model adaptation or multi-threading. However, the resource contention by multiple threads/apps may lead to performance degradation, and the

real-time model adaptation can cause a significant accuracy drop. More importantly, their requirement of full use of hardware resources for RTOD can be problematic when applied to a real-time multi-tasking platform.

1) *Why cannot Heo et al. [9] be modified for multi-neural network execution?*: The authors of [9] showed that the WCET model for a certain layer can work in multi-tenant systems, but they neither showed how to share the limited resources between concurrent DNN models nor accounted for the overhead of context switching between apps. In contrast, DynaMIX addresses these problems in multi-tasking systems.

C. Hardware-based Scheduling for Multi-DNN Inferences

AI-MT [11] proposed a new NN accelerator architecture and scheduling scheme for multi-tasking platforms. [12], [13] coordinated multiple latency-critical DNN tasks by using specific DNN accelerators (e.g., FPGA, NPU). Scheduling jobs across heterogeneous processors allowed their real-time execution. The most recent work in [14] proposed a GPU scheduling method to run multiple real-time apps. Basically, it schedules memory- and computation-bound jobs across heterogeneous or several processors to enhance concurrency. However, it does not work in the representative case of Fig. 1. It only focuses on scheduling multiple real-time jobs with the unmodified DNN models which require a *static* amount of resources, thus limiting the number of concurrent jobs it can handle. In contrast, DynaMIX dynamically adjusts the resource consumption of each app. Moreover, DynaMIX is orthogonal to these prior solutions, and hence can run with them together.

III. BACKGROUND: MIXED-PRECISION QUANTIZATION

Quantization of activation and model parameters can accelerate DNN execution by reducing the computational complexity of the underlying models. This is predicated on the fact that integer operations yield a much higher throughput in vectorized computations than floating-point operations on most computing platforms. However, removing some bits in a fully-trained model (that is already converged to the lowest loss) causes an output perturbation between the full-precision and quantized models, thus degrading accuracy significantly.

Mixed-precision quantization has been explored as a promising solution to this problem by using layers of different bit-widths [21]–[25]. Using a higher bit-width at a layer more sensitive to quantization can help the layer preserve its original values, thus making the model suffer less output perturbation and accuracy drop. Typical mixed-precision quantization algorithms are composed of layer sensitivity measurement and layer bit-width decision-making.

A. Measurement of Layer Sensitivity

Layer sensitivity represents the extent to which the model output changes when a certain layer is quantized. For the most exhaustive mixed-precision approach, a model of L layers with B types of bit-widths yields B^L -quantized models. Since DNN models are recently becoming deeper, mixed-precision will be less attractive. So, identifying layer sensitivity is an efficient

way to reduce the large design space for bit allocation. Well-measured layer sensitivity can also be used to calculate the overall perturbation of the resulting mixed-precision model, and is a good measure for finding the best bit setting.

Layer sensitivity is affected by several factors, such as layer position, operation type, connection with other layers, and layer parameter size. It is difficult to define sensitivity by considering all these factors in a large model; researchers used output perturbation as their sensitivity metric by calculating $L2\text{-Norm}$ or $KL\text{-Divergence}$ [22], [25].

In this paper, we use the sensitivity metric defined in the state-of-the-art mixed-precision quantization method ZeroQ [22], which is based on $KL\text{-Divergence}$ between the full-precision model and the quantized model, as:

$$S_i(k) = \frac{1}{N} \sum_{j=1}^N \text{KL}(\mathcal{N}(x_j), \tilde{\mathcal{N}}_i^k(x_j))$$

where $S_i(k)$ denotes the sensitivity of quantized model $\tilde{\mathcal{N}}_i^k(\cdot)$ in which the i -th layer is quantized into k -bits, $\mathcal{N}(\cdot)$ indicates the full-precision model, and x is a small set of input images of size N used for sensitivity measurement.

B. Hardware-aware Bit-width Decision

Mixed-precision quantization has the flexibility of hardware-aware model compression. As the most basic method of bit-width decision, rule-based schemes have been used based on the knowledge of the DNN model and hardware architecture as well as manual effort. However, increasingly complex models make such heuristics difficult to apply. ZeroQ [22] employed a *Pareto frontier*-based method that finds an optimal compressed model with minimum output perturbation. The overall sensitivity for each mixed-precision model is computed by summing the sensitivities of all layers in the model. Although the authors did not reflect inter-layer dependency in this process, they showed that such sensitivity calculation incurs less computational overhead and produces good empirical results. Additionally, sensitivity and model size are considered in ZeroQ as the indirect indicators of accuracy and latency.

However, each processor has a distinct architectural design, implying that the best bit settings vary with hardware [21]. For instance, a weight parameter layout designed to increase reusability can efficiently reduce the latency of conventional convolution layer, which can thus have a higher bit-width than fully-connected or depth-wise convolution layers. Thus, with hardware-software codesign, this advanced quantization can reduce the computation and resource costs without any severe loss of inference accuracy compared to hardware-oblivious DNN compression methods (e.g., uniform quantization [15]).

Thanks to its flexibility, mixed-precision can accelerate the network model with a negligible accuracy degradation. Thus, recent processing engines have been released with the function of mixed-precision arithmetic with variable bit-widths [26], [27]. These recent advancements have raised the importance of mixed-precision-based methods for DNN acceleration.




Fig. 2. Runtime model reconfiguration and execution mechanism.

Inspired by the reinforcement learning-based bit-width setting method [21], we determine the best mixed-precision model for each app in a hardware-aware manner. Instead of using indirect indicators, e.g., FLOPs, model size, and sensitivity, we use the worst-case latency (to meet the deadline), peak memory usage (to consider memory capacity), and inference accuracy measured in our simulated environment. We aim to use a mixed-precision approach to find an optimal solution that maximizes the total accuracy while making the most of limited resources. Sec. IV and V will detail how DynaMIX builds the mixed-precision model and chooses an optimal bit configuration for each DNN-based app.

IV. DYNAMIC MODEL RECONFIGURATION & EXECUTION

At the core of DynaMIX is finding the most suitable mixed-precision model for each app by reflecting the current system condition; app processes may otherwise be unstable for DNN models due to their low accuracies or severely underutilize the platform resources. Thus, we need to consider various compressed models with different bit configurations, their saving and loading of many models. However, reloading the resulting entire model each time may be neither suitable for real-time apps nor efficient when only a few layers need to be loaded. Thus, we propose to reconfigure the model bit-width configuration at runtime by lazy loading FP layers in order to reduce the model loading overhead and ensure deadline satisfaction. This dynamic model reconfiguration approach is inspired by [28], [29] where bit-width decisions are made only for a *single* DNN model, and training a special DNN from scratch is necessary for runtime adaptive quantization. In contrast, we develop a solution that determines bit configurations for multiple DNNs without requiring any additional (re)training.

In DynaMIX, when the optimal bit setting for each app is determined, only the additional layers for the target model are loaded to compose the path (⑦ in Table II). Before operating this at runtime, we separate a model into layers and store the full-precision (FP) and quantized (INT) versions of all layers (①). That is, $2L$ versions of layers are saved for a model consisting of L layers. Specifically, operation type (e.g., convolution) and FP weights are saved for an FP layer; for an INT layer, additional parameters required for quantization of FP activations are saved with operation type and INT weights.

TABLE II
NAVIGATION FOR SYSTEM COMPONENTS. PIC MEANS THE PROCESS IN CHARGE OF THE CORRESPONDING COMPONENT.

Stage	Component	Section	PIC
Offline	① Per-precision Layer Saving	IV	Sub
	② Mixed-precision Model Building	V-A	Sub
	③ Profile Generation	V-B	Main
	④ Lookup Table Generation	V-D	Main
Online	⑤ Memory Resource Allocation	V-C	Main
	⑥ Lookup Table Querying	V-D	Main
	⑦ Model Reconfig. and Execution	IV	Sub

Mechanism. Model reconfiguration consists of 1) loading necessary layers and 2) changing the model execution path. Here we deal with a multi-tasking system in which the main process is in charge of process monitoring and memory allocation, and the subprocess is in charge of model execution. To load only the layers needed at runtime, DynaMIX should first hold the fully-quantized models of all possible apps as their base models by loading their INT layers and building them without changing the original architecture. When a new camera image enters to the system and any change in the status of concurrent apps is detected, the main process adjusts the amount of memory resource allocated to the concurrent apps (⑤). After determining the new bit setting from the adjusted memory space for each app, the necessary FP layers required for the new model are loaded (⑥), and then they are executed instead of the existing INT layers. After execution, the FP layers are deleted, and each model returns to its original base model state. Fig. 2 shows the above mechanism with two different cases. How to build a mixed-precision model while considering the given memory space will be elaborated next.

V. SYSTEM ARCHITECTURE

Our main objective is to optimize the allocation of system resources, such as CPU time and memory space, for multiple simultaneous apps while meeting the deadline, which has not yet been addressed adequately because of the high computational complexity of real-time CNN models.

Representing a model using both high and low precisions (②) is key to reduction of the computational cost of DNNs, thereby enabling multiple DNN apps to share limited platform resources. Central to DynaMIX is a memory resource allocation algorithm that determines the best bit setting for each app to use so as to maximize total accuracy (⑤). This algorithm uses app profiles that are created offline to save the information of a set of mixed-precision models (③), as a constraint term. Thus, optimization of computation resource is triggered by the apps whose memory requirements change dynamically. After completing a series of processes for resource optimization, the additional layers to load/execute are selected by looking up tables (⑥), which were generated offline (④).

To reduce the online burden of computation-heavy processes, DynaMIX is divided into offline and online stages (see Fig. 3). ①, ⑦ are discussed in Sec. IV, and the remaining components are elaborated in the subsequent sections.

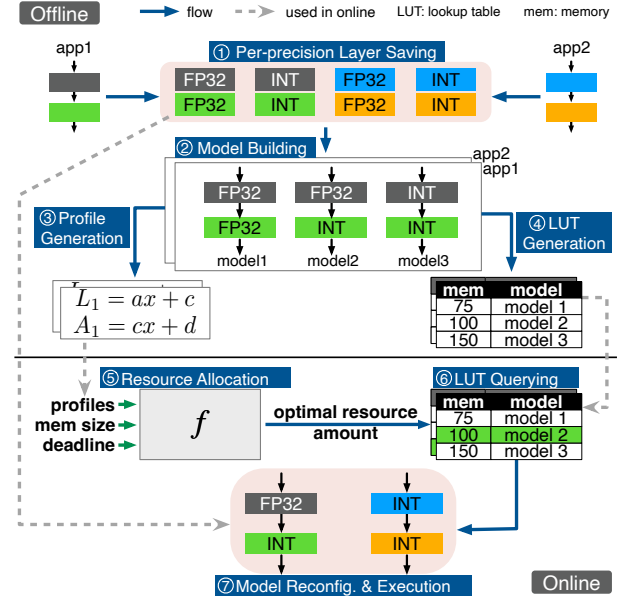


Fig. 3. System architecture.

A. Mixed-Precision Model Building

Using the FP and INT versions of all layers, we can build different mixed-precision models (based on the identical architecture) which show the effects of mixed-precision quantization in terms of accuracy, speed, model size, etc. The main objective of DynaMIX is to find a model that maximizes the accuracy with the given resources and deadline. We generate different compressed models of various bit configurations (showing different performances), and use them for the prediction of performance in accordance with the amount of available memory resource.

We develop a sorting-based approach to build models while yielding the highest accuracy compared to the degree of compression (defined as the number of the quantized layers in the model). In this approach, the FP layers of the model are sorted in descending order of layer sensitivity discussed in Sec. III-A. ZeroQ [22] shows the overall sensitivity of a quantized model (where multiple layers are quantized with low bit-widths) can be calculated by summing up the sensitivities of all quantized layers in that model. Based on this calculation, starting with the fully-quantized model, the INT layers are replaced with their corresponding FP layers one-by-one (in descending order of layer sensitivity). This means that at all degrees of quantization, the generated model would yield the smallest output perturbation from the full-precision model, thus leading to the smallest accuracy drop. This model building process is described in Algorithm 1. Although this can be presented recursively or iteratively like depth-first search, we present it recursively for readability, and the initial invocation for this procedure is BUILD (Fully-quantized model, \emptyset).

Our sorting-based approach yields L compressed models for a model consisting of L layers, which effectively removes unnecessary and repetitive models among 2^L models produced by an exhaustive model generation approach.

Algorithm 1 Mixed-Precision Model Building

L : a set of layers sorted in descending order by layer sensitivity

S : a set of compressed models ($S \leftarrow \emptyset$)

procedure BUILD(model m , layer l)

$n = \text{Replace } l \text{ of } m \text{ with FP layer}$ \triangleright If l is \emptyset , $n \leftarrow m$

$S \leftarrow S \cup n$

v : the most sensitive layer in L \triangleright If L is \emptyset , **return** S

remove v in L

BUILD(n , v)

return S

end procedure

Fig. 4 shows the inference results of the compressed models of VGG16 [30], ResNet50 [31], and YOLO-v3 [20]. Note latency includes the time for model reconfiguration, execution, and restoration. We used 2 bit-widths, FP32 and INT8, throughout this paper to demonstrate the effectiveness of using mixed-precision models on app performance, i.e., all layers are expressed with either precision. The three graphs in the upper row show latency changes and the three graphs in the lower row display accuracy changes as a function of the degree of compression. Considering memory usage and latency of a model may change each time we run the model, we measured 100 times for each model and used the highest peak memory usage and latency as the worst case. The results show a latency reduction without sacrificing accuracy much when compression degree is increased. To utilize the performance results of the compressed models in the online stage, two profiles and a lookup table are created for each model.

B. Profile Generation

The above process yields a number of mixed-precision models for each app, and hence the goal of our memory resource allocation is to find an optimal set of compressed models for all concurrent apps. Although such a large number of combinations offers flexibility to accommodate various real-world cases, the high computational cost of finding an optimal set makes it unsuitable for real-time resource allocation. Our model compression yields $O(L)$ models, and in the case of a 10 FPS camera, the end-to-end latency per frame should be less than 100 ms [3]. Given that I apps are simultaneously running on the platform, an excessive amount of time (up to $O(100 \cdot L^I)$ ms) is required to find an optimal combination.

For real-time optimization, the authors of NestDNN [7] considered only five candidate models, which come at the cost of both resource utilization and accuracy. In DynaMIX, however, profiles are generated offline for each app in the form of polynomials for latency and accuracy by applying polynomial regression to the compressed models. This app performance profiling eliminates the need of additional inference processing in the online stage. We introduce latency and accuracy profiles, which are used in our online memory resource allocation algorithm.

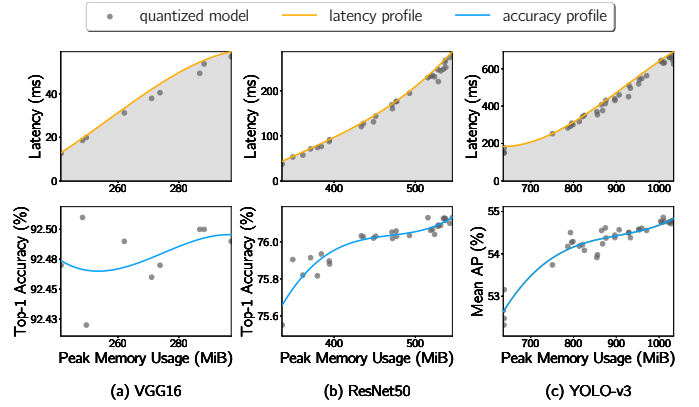


Fig. 4. The results of model compression and profiling on VGG16, ResNet50, and YOLO-v3. Each gray dot indicates a certain compressed model that has a specific bit setting. A yellow line denotes the profile of worst-case latency, thus covering all the compressed models as shown in the gray areas. A blue line means the profile of most-likely accuracy, showing the trend of accuracy.

TABLE III
LATENCY AND ACCURACY PROFILES. X MEANS PEAK MEMORY USAGE.

NN Model	Type	Profile
VGG16	Latency	$-0.0001472 x^3 + 0.1135 x^2 - 28.16 x + 2267$
	Accuracy	$-7.781e-07 x^3 + 0.0006417 x^2 - 0.1753 x + 108.4$
ResNet50	Latency	$1.021e-05 x^3 - 0.01083 x^2 + 4.648 x - 683$
	Accuracy	$1.285e-07 x^3 - 0.000181 x^2 + 0.08553 x + 62.48$
YOLO-v3	Latency	$-6.706e-06 x^3 + 0.01861 x^2 - 15.5 x + 4241$
	Accuracy	$6.764e-08 x^3 - 0.0001834 x^2 + 0.1676 x + 2.821$

1) *Latency Profile*: By employing the latency profile of a target DNN, we can estimate the worst-case latency when the model is compressed to a certain size. If some models cannot be covered by the latency profile, the newly selected model at runtime can exceed the expected number obtained from the latency profile, thus failing to guarantee the deadline. Thus, the latency profile should cover all the resulting compressed models. The impact of inter-task interference in a multi-tasking platform on execution latency is discussed in Sec. VI-B.

To create the latency profile that covers all compressed models, the compressed models are stored when its latency is greater than the latencies of the smaller models, and then polynomial regression is applied to the stored models. We exploit cubic polynomials because latency is not completely but fairly proportional to peak memory usage, and high-order regression burdens the optimization algorithm (see Fig. 4). After obtaining a polynomial profile for each model, to guarantee the worst-case latency for a given memory space, we delete those models whose latencies are greater than the values estimated with the latency profile from a set of all compressed models. The resulting latency profiles of VGG16, ResNet50, and YOLO-v3 are provided in Table III.

2) *Accuracy Profile*: The role of the accuracy profile is different from that of the latency profile where we observe the change in accuracy according to the degree of compression. For this, we apply regression to all of the remaining models after creating the latency profile, without deleting models any further. Accuracy also has a weak positive linear correlation with peak memory usage, and thus regression is performed using a cubic polynomial (see Fig. 4). The accuracy profiles of VGG16, ResNet50, and YOLO-v3 are provided in Table III.

C. Optimal Resource Allocation

Once a new camera image enters to the system and any change in the status of app processes is detected, DynaMIX allocates the memory resource for concurrent apps while maximizing the total accuracy under the given deadline. The optimization problem in this process is solved by reflecting the latency and accuracy profiles generated in the offline stage to determine the optimal allocation of memory resource for each app. In what follows, we present the development of a resource allocation algorithm, which is formulated as a constrained nonlinear optimization problem.

DynaMIX assumes a limited memory capacity, and the timing requirement of a real-time app is determined by camera's frame rate. Therefore, the problem of resource allocation for concurrent apps is formulated as:

$$\begin{aligned}
& \arg \max_{m_i} \sum_{i \in I} \lambda_i A_i(m_i) \\
& \text{s.t.} \quad \sum_{i \in I} L_i(m_i) \leq D - \varepsilon \\
& \quad \forall i \in I : m_i \leq M_{\max} - \mu_i \\
& \quad \forall i \in I : m_i^L \leq m_i \leq m_i^U
\end{aligned} \tag{1}$$

where I is the set of concurrent apps. We want to find the optimal memory size m_i for each app $i \in I$ that maximizes total accuracy. *Total accuracy* is calculated as the weighted sum of accuracies of all apps, and λ_i is used to give different weights when each app has a different level of importance or the degree of change in accuracy during quantization varies with app. λ_i is set to 1 in this paper. For the constraint terms, we consider memory capacity apps can use M_{\max} , and the given deadline D . Here, ε is the average processing time in the main process, and μ_i is the memory usage of the base models of concurrent apps, which is calculated as $\mu_{total} - \mu_{app_i}$, where μ_{total} is $\mu_{vgg} + \mu_{resnet} + \mu_{yolo}$, and μ_{app_i} is one of μ_{vgg} , μ_{resnet} , and μ_{yolo} . ε and μ_i will be detailed in Sec. VI-B in which we assess the feasibility of solving this problem. We also use the latency profile L_i and accuracy profile A_i obtained from the profile-generation phase. The latency-related constraint term ensures that estimated latency $L_i(m_i)$ does not exceed the deadline. Lastly, we set lower and upper bounds (m_i^L, m_i^U) for m_i to reduce the search space in solving this optimization problem. m_i^L is set to the smallest peak memory usage of app i , and m_i^U is set to the largest peak memory usage of app i . The optimization result determines the appropriate memory space size to be allotted to each app.

Algorithm 2 Lookup Table Generation for App i

```

 $\mathcal{T}$  : lookup table ( $\mathcal{T} \leftarrow \emptyset$ )
 $\mathcal{M}$  : a set of mixed-precision models
 $z$  : The number of lookup table entries
 $max(i)$  : the largest peak memory usage
 $min(i)$  : the smallest peak memory usage
 $m_{pm}$  : peak memory usage of model  $m$ 
 $m_{acc}$  : accuracy of model  $m$ 

Sort  $\mathcal{M}$  in ascending order by peak memory usage
 $val \leftarrow 0, oldkey \leftarrow 0$ 
while  $m \in \mathcal{M}$  do
     $key = min(i) + \left\lceil \frac{(z-1)(m_{pm}-min(i))}{max(i)-min(i)} \right\rceil \frac{max(i)-min(i)}{(z-1)}$ 
    if  $key == oldkey$  then
        if  $m_{acc}$  is greater than  $val_{acc}$  then
             $val \leftarrow m$ 
        end if
    else
         $\mathcal{T} \leftarrow \mathcal{T} \cup (oldkey, val)$ 
         $val \leftarrow m, oldkey \leftarrow key$ 
    end if
    remove  $m$  in  $\mathcal{M}$ 
end while

```

D. Lookup Table for Final Model Selection

After resource allocation, DynaMIX should determine which compressed models to load into the system. Given a large number of compressed models for each DNN model, searching for the best fitting model sequentially will likely take too long for real-time apps. We thus reduce this search time via lookup tables, which are generated for each model in the offline stage and used in the online stage.

1) *Generation of Lookup Table*: We present a simple yet efficient method for generating the lookup table, which is based on the principle of arithmetic sequences. We first divide equally the range of peak memory usage of each app. Second, a new entry is added to the table in the form of a key-value pair, where the key is the endpoint of a certain divided section, and the value is the bit setting of the model with the highest accuracy when the memory size corresponding to that end-point is allocated. Starting from the smallest section, the entry for each section is iteratively inserted into the table. As a result, the resulting lookup tables remain to be of constant size (i.e., the preset number of table entries) regardless of the complexity of the DNN model. Fig. 5 shows the accuracy of the models of which bit settings are stored in the lookup table. Each graph shows the number of different settings DynaMIX needs to store for each app (e.g., 2 for VGG16). The detailed procedure for lookup table generation is described in Algorithm 2.

2) *Querying Lookup Tables*: The generated lookup tables allow DynaMIX to select the final model quickly after the online resource allocation.




Fig. 5. The accuracies of the models stored in the lookup table. The ticks marked on each x-axis denote the sections created by dividing the range of peak memory usage by a fixed number (i.e., the number of lookup table entries). Each of the resulting lookup tables contains 50 entries.

$$key = \min(i) + \left\lfloor \frac{(z-1)(m_i^* - \min(i))}{\max(i) - \min(i)} \right\rfloor \frac{\max(i) - \min(i)}{(z-1)}$$

where m_i^* denotes the memory size allocated for app i , and z is the number of table entries. $\max(i)$ and $\min(i)$ represent the maximum and minimum peak memory usages of app i .

After entering the resulting keys into the tables, the final models to execute are determined. Each subprocess of an app then loads the additional FP layers required for the newly selected model and runs according to the reconfigured execution path as discussed in Sec. IV.

VI. EVALUATION

We now evaluate DynaMIX by implementing all the components in Table II in a simulated environment. We implemented layer quantization and model execution using PyTorch 1.10.2 [32], which is one of the most representative deep learning frameworks. In compliance with the quantization support by PyTorch, INT8 is used as low precision on x86-64 CPU for the acceleration of DNN execution. Instead of using indirect metrics for app performance, we measured actual latency by running all the compressed models and peak memory usage by using memory profiler [33]. For profile generation, we used the polynomial regression utility, polyfit, in the NumPy API. The resource allocation algorithm was programmed using the sequential least square programming utility in SciPy API. All our experiments were conducted on an AMD Ryzen 7 5700G processor and 16GB RAM with Ubuntu Linux 20.04 LTS operating system. Given the goal of DynaMIX, we want to meet the same deadline of concurrent DNN apps.

Following [9], [10], the target deadline is relatively determined by reflecting the performance of the original (i.e., full-precision) models measured in our environment. Even though we experimented in a simulated environment, we aim to show DynaMIX's ability of resource optimization when the state of the running apps changes.

Multi-tasking Implementation. Our multi-tasking environment consists of the main and sub-processes. The main process detects changes in the system, and allocates platform resources to concurrent apps. Each subprocess is responsible for the execution of its DNN model. Note that the number of subprocesses is equal to that of types of possible apps. For multi-tasking, we use the Non-Preemptive Shortest Job First (SJF) scheduling to avoid the high overhead of context switching between DNN-based apps, consider the equal priority between apps (attributed to the common deadline), and minimize the

average waiting time of concurrent apps. For example, the three apps introduced throughout this paper are scheduled in the order of VGG16, ResNet50, and YOLO-v3.

As discussed in Sec. IV, each subprocess first builds the fully-quantized model of its app as a base model. When processing images coming from the camera, DynaMIX behaves differently depending on whether any change in the status of concurrent apps is detected or not. In case a change occurs, the main process determines the optimal memory space for each app, and then notifies the corresponding bit configuration to one of the subprocesses by sending a message (*case 1*). Otherwise, the main process does not make resource allocation and notifies the last bit configuration to the subprocess (*case 2*). When the subprocess receives the information, it loads the necessary layers, runs the model, and then reports to the main process that the job is over through shared memory. Then, the main process repeats this procedure until all apps are executed. In the experiments in Sec. VI-C, we adjusted the status of apps at 1s intervals. If *case 1* meets the deadline, so does *case 2*. We therefore show the results of deadline satisfaction for *case 1*.

In what follows, we first present the DNN models and datasets used in our evaluation, and discuss the feasibility, effectiveness, and robustness of DynaMIX.

A. Neural Networks and Datasets

Image classification and object detection apps often run simultaneously on AVs. Furthermore, even for identical tasks, the complexity of model varies greatly with the function. For example, classifying objects in front of a car is more difficult than classifying the color of a traffic light. Considering the diverse real-world cases, the following three DNN models were selected for two types of apps. Although the state-of-the-art models change rapidly, their underlying architectures are predicated on the models of our choice [34], [35].

1) *VGG16 on CIFAR-10*: CIFAR-10 [36] is made up of 60K images in 10 classes, consisting of a low-resolution color image of animals or vehicles. Classification of CIFAR-10 is considered to be a simple and easy problem owing to the small number of classes and the clear visual distinction between class objects. Based on CIFAR-10, we used VGG16,¹ which is fully trained on the dataset. VGG is one of the most popular and representative architectures [7], [24], [25], which uses small (e.g., 3×3) kernels in all convolution layers. In particular, VGG16 (consisting of 16 layers) shows nearly state-of-the-art inference accuracy on the CIFAR-10 dataset. As a result of the simplicity of the task, VGG16 shows higher accuracy than the other models as shown in Figs. 4 and 5.

2) *ResNet50 on ImageNet*: ImageNet [37] is composed of 1.4M images in 1K classes, each of which is a high-resolution color image of animals, plants, vehicles, or electronic devices. Classification on ImageNet is more difficult than on CIFAR-10 because of the high similarity of features between classes and a large number of classes. The ResNet50 model,² fully trained

¹Obtained from <https://github.com/chengyangfu/pytorch-vgg-cifar10>

²Obtained from <https://github.com/pytorch/vision/blob/main/torchvision/>




Fig. 6. Distributions of the latency and peak memory usage of main and sub-process. If the number of subprocesses is 1, YOLO-v3 is loaded and executed. In case of 2, ResNet50 and YOLO-v3 are loaded and executed, and in case of 3, all three apps are loaded and executed. In the last graph, ‘V’, ‘R’, and ‘Y’ stand for ‘VGG16’, ‘ResNet50’, and ‘YOLO-v3’, respectively. Each box contains the results of 100 measurements.

on the ImageNet dataset, was used in this study. ResNet, released after VGG, has a much deeper neural network architecture (a depth of up to 152 layers) than VGG (a depth of up to 19 layers), which solves the vanishing gradient problem caused when designing deep networks by adding skip connections to the network. Due to the complexity of ImageNet, the accuracy of ResNet-50 is lower than that of VGG16 on CIFAR-10 (see Figs. 4 and 5), but currently widely-used ImageNet-based models show accuracies of 75 ~ 80%, and ResNet-50 shows nearly state-of-the-art inference accuracy on ImageNet. More importantly, the ResNet architecture has great potential for building the current high-performance deep model [38].

3) *YOLO-v3 on MS COCO*: For the object detection task, we used the YOLO-v3 model³ trained on the MS COCO dataset [39], which is composed of 5K images in 91 classes, consisting of images pertaining to our daily life such as one or more objects of vehicles and daily necessities. YOLO is the first object detection architecture that localizes and classifies multiple objects in an input image or video in one single forward propagation. Thanks to its speed, RTOD has been realized in autonomous driving [3], and its role in promoting safety and convenience continues to grow. In Figs. 4 and 5, the accuracy of YOLO-v3 (measured in “average precision (AP)”) is lower than that of the classification models (measured in “Top-1”). This is because AP means not only the accuracy of classification but also the prediction accuracy of each object position, whereas Top-1 only indicates classification accuracy. However, YOLO-v3 performs well in balancing accuracy and latency and is one of the most commonly used RTOD models in both academia and industry [3], [4], [40].

B. Feasibility of Resource Allocation Algorithm in DynaMIX

Ensuring the feasibility of resource allocation depends on the following factors: 1) the overhead and additional memory usage in solving Eq. (1), 2) the reasonableness of using latency profiles for the first constraint term of Eq. (1), and 3) the memory usage of each app’s base model for the second constraint term of Eq. (1). We show the feasibility of the proposed algorithm by identifying each of these factors.

1) *Runtime overhead and memory usage in solving Eq. (1)*: Eq. (1) entails a cubic objective function with cubic and linear constraints. Although the size of the search space is reduced by using lower/upper bounds, we need to assess its runtime

overhead to ensure the deadline satisfaction. Fig. 6 (a) shows the number of concurrent apps vs. the latency of solving Eq. (1). Latency in the single-tasking case was about $1.7\times$ smaller than the multi-tasking case. Based on this observation, we set ε in Eq. (1) to 15 ms by considering the maximum latency of resource allocation. Meanwhile, Fig. 6 (b) shows the peak memory usage is found almost constant regardless of the number of running apps (the ratio of the smallest to the largest mean is approximately 1). So, given that M_{max} is the memory capacity for the subprocesses of the apps, not for the main process, the main process does not require more space and hence has little effect on M_{max} .

For the various cases in real-world, the latency and peak memory usage values (for each box) in Figs. 6 (a) and (b) are measured while changing deadline and memory capacity.

2) *Reasonableness of using latency profiles*: Given the latency profile is based on the measurements in isolation, the latencies should hold even when multiplexed with other apps. From Figs. 6 (c) and (d), one can observe that the latency of each compressed model holds in the non-preemptive SJF scheduling-based multi-tasking environment. Due to space limitation, we only showed two models: (c) for the fully-quantized model of YOLO-v3, and (d) for its full-precision model. From these, we conclude that the latency profile can be used for the first constraint term.

3) *Memory usage of each app’s base model*: More concurrent apps occupy more memory space by loading their base models. Fig. 6 (e) shows the memory used by each base model. Based on their mean values, we set μ_{vgg} , μ_{resnet} , μ_{yolo} , μ_{total} to 228, 273, 323, 823, respectively, in the second constraint term.

C. Effectiveness of DynaMIX

To show the efficacy of DynaMIX in a multi-tasking environment, we designed scenarios where the number of concurrent apps changes under two types of timing constraints: sufficient and tightened deadlines to run multiple DNN apps (see Table IV). Owing to the space limitation, we selected some deterministic yet representative (instead of random) cases where deadline or the number of concurrent apps is adjusted to show the change of each app performance.

In general, the model execution speed depends strongly on the processing engine (e.g., CPU, GPU). Considering the use of much slower CPU than GPU, we set the deadlines based on the latencies measured in our simulated environment. In

³Obtained from <https://github.com/eriklindernoren/PyTorch-YOLOv3>




Fig. 7. Resource optimization results. In (a), we marked the margin between peak memory usage by 3 apps and the memory capacity with a blue line; ‘V’ is VGG16, ‘VR’ is the sum of memory used by VGG16 and ResNet50, and ‘VRY’ denotes the total memory space used by all 3 apps. (b) shows the total latency (consisting of model reconfiguration, app execution, etc.) and each deadline (marked in a blue line), and (c) shows the accuracy change of YOLO-v3.

TABLE IV
EVALUATION SCENARIOS ON A MULTI-TASKING SYSTEM

Scenario	Description
Initial	YOLO-v3-based app is working alone (0).
1	VGG16-based app starts to run in (1).
2	Deadline decreases from 700 ms to 500 ms in (1).
3	Deadline decreases from 500 ms to 250 ms in (2).
4	ResNet50-based app starts to run in (0).
5	Deadline decreases from 700 ms to 500 ms in (4).
6	Deadline decreases from 500 ms to 250 ms in (5).
7	VGG16-based app starts in (4).
8	Deadline decreases from 700 ms to 500 ms in (7).
9	Deadline decreases from 500 ms to 250 ms in (8).

particular, we used minimum (148 ms), median (470 ms), and maximum (702 ms) worst-case latencies resulting from YOLO-v3’s mixed-precision models. Memory capacity was set to 1600 MiB by considering the maximum peak memory usage of those models. Figs. 7 (a) and (b) show that DynaMIX’s resource optimization enables deadline satisfaction.

We assume the initial environment (0) to be in a state where YOLO-v3 operates alone and there is enough time and memory space to run the full-precision model. In such a case, DynaMIX assigns 1026 MiB to the app, expecting the highest accuracy when the accuracy profile is used. Although the target deadline (700 ms) is determined while considering the full-precision model, DynaMIX selects a different model (INT layer accounts for 21% of the total) which runs more quickly and accurately with less resources. In (1), despite sharing resources with VGG16, YOLO-v3 uses the same bit configuration used in (0). This phenomenon is attributed to the low memory usage and time complexity of VGG16. Also, given the accuracy of VGG16 is almost constant regardless of the degree of compression (see Fig. 4), DynaMIX compresses VGG16 maximally and allocates most resources to YOLO-v3. However, in (4), the execution of ResNet50 requires compression of both ResNet50 and YOLO-v3. Specifically, the memory space allocated to YOLO-v3 decreases from 1017 MiB to 952 MiB, and hence YOLO-v3 uses more low precision layers (60% of the total) than (1). In (7), the degrees of compression of ResNet50 and YOLO-v3 are the same as in (4) due to the low computation and memory costs of VGG16.

(2), (5), and (8) show DynaMIX’s resource optimization for concurrent apps when the deadlines are tightened (from

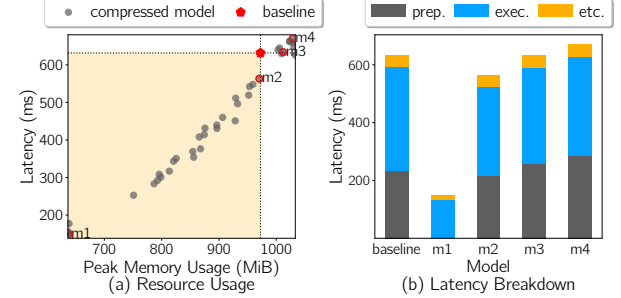


Fig. 8. Comparison of results when DynaMIX and the baseline are applied to YOLO-v3. In (a), the red pentagon represents the worst-case latency and peak memory usage when using the baseline, each of which is calculated using 100 measurements. (b) shows the latency breakdown of the main models in (a). In the case of the baseline, “prep.” means the time required in loading the entire model; otherwise, it means the time for model reconfiguration. “etc.” includes the time spent in model restoration, interprocess communication.

700 ms to 500 ms). (3), (6), and (9) show the optimization results with the deadlines tightened further (from 500 ms to 250 ms). In these scenarios, concurrent DNN models are executed by building more compact mixed-precision models. Consequently, a larger margin between the peak memory usage and memory capacity is made available in Fig. 7 (a). The overall margins between the measured latency/peak memory usage and deadline/memory capacity may seem large, because we used the worst-case latency and maximum peak memory usage (through iterative measurements) when profiling, and hence the margins vary with the execution.

Our resource allocation aims to maximize total accuracy without violating any timing constraint, and hence there is no significant loss in accuracy (see Fig. 7 (c)). Due to the lack of space, we presented the results of YOLO-v3 only, the most influential model for both real-time execution and total accuracy drop. In all the scenarios considered, all three models show less than 2.2% accuracy drops. The maximum accuracy drop occurs in YOLO-v3 in (9).

D. Performance Comparison

1) *Comparison with baseline:* If a new app is allotted enough resources to run the full-precision model in **case 1**, the subprocess will reconfigure the DNN pipeline by loading all its FP layers, which may use resources less efficiently than the baseline approach (i.e., loading the entire FP model at once

without reconfiguration). Albeit such inefficiency, we show DynaMIX to make efficient use of resources in most cases.

Fig. 8 (a) shows the peak memory usage and worst-case latency of the mixed-precision models of YOLO-v3 (the same as the YOLO-v3 latency graph in Fig. 4), but this time, the peak memory usage and worst-case latency resulting from the baseline approach are marked with red pentagon, and the latency breakdowns for the main cases (marked with red hollow dots) are shown in Fig. 8 (b). This analysis suggests several aspects of DynaMIX’s effectiveness. First, *m4* (all layers in the execution path are replaced with FP layers) shows inefficient latency and memory usage compared to the baseline. However, 73.7% of our compressed models showed less usage of memory and processing time (the yellow area in Fig. 8), especially 78.9% of our compressed models showing reduced latencies. In case of *m1* (all layers are in low precision), latency and memory usage were reduced by $4.3\times$ and $1.5\times$, respectively, compared to the baseline. Additionally, when about 22.7% of the layers were quantized (*m3*), it showed the same level of latency as the baseline, and when about 30.7% of the layers were quantized (*m2*), it showed the same level of memory usage. From these results, we conclude DynaMIX reacts well to the given deadlines by optimizing the computation and memory resources required in running apps.

2) *Comparison with NestDNN*: NestDNN [7] also uses a DNN pruning method to run concurrent vision apps (non-real-time). They create a fixed number of small models according to the degree of pruning. They demonstrated feasibility in terms of memory usage, accuracy, and energy consumption. Energy consumption is not in our scope and the app’s memory usage is discussed by using its model size as an indirect metric. So, we compare DynaMIX with NestDNN only for accuracy.

Table V shows the differences between the generated models when the two methods are applied to VGG16. Here, DynaMIX is demonstrated: 1) for all the generated mixed-precision models (ALL) and 2) for the models saved in the lookup table (LUT). While NestDNN uses only five compressed models for online resource allocation, we can handle more models by inserting the app profiles into the optimization problem and reducing the time of searching for models via lookup tables. Furthermore, all of the compressed models generated by DynaMIX show higher accuracy than NestDNN, because its pruning scheme can incur significant accuracy drop since the distribution of model parameters is changed. Although the authors recover the accuracy to some extent via retraining, the accuracy is not fully recovered due to the diminished capacity of the model. The higher accuracy of DynaMIX enhances the feasibility, but shows a lower compression ratio. However, the degree of compression can be increased by using lower precision (e.g., INT4), and if different low precisions are used simultaneously, we can increase both the compression degree and accuracy.

E. Robustness of DynaMIX

Since system error in AVs can cause exigent situations, we must ensure system robustness in terms of system availability

TABLE V
COMPARISON BETWEEN NESTDNN [7] AND DynaMIX. MAX COMPRESSION RATIO MEANS THE RATIO OF THE UNCOMPRESSED MODEL SIZE TO THE SMALLEST COMPRESSED MODEL SIZE.

	[7]	DynaMIX (All)	DynaMIX (LUT)
# Compressed models	5	17	2
Min accuracy	83	92.47	92.47
Max accuracy	89	92.51	92.51
Max compression ratio	7.8	4.0	4.0




Fig. 9. Total accuracy while changing the deadline and memory capacity. The lighter the color, the higher the total accuracy. Black cells denote invalid results caused by the inability of optimization or a significant drop in accuracy.

and quality — two critical factors of system failure. The resource optimization mechanism in DynaMIX fails to yield results if the deadline is set below a certain level. Furthermore, heavily compressed models can be used to meet the short deadline but may lead to fatal accidents due to their low accuracy. Considering these factors that affect system availability and quality, we can evaluate system robustness.

Below we show the robustness of DynaMIX using the total accuracy resulting from lookup tables. Since the robustness may be different in the platforms with different memory capacity, we visualized the relationship between deadline, memory capacity, and total accuracy in Fig. 9. The black cells indicate the points where system gives invalid results for three reasons: when available bit-widths cannot meet the deadline, when memory capacity cannot support DynaMIX, or when the total accuracy is degraded by more than a predefined threshold (3% in this paper). In the remaining cases, DynaMIX guarantees acceptable quality for users.

As our environment supports INT8 and FP32 operations, total accuracy does not decrease by more than 3%, but the lower bounds of deadline and memory capacity appear high when the black area is considered. The ultra-low-bit integers, such as INT4 and INT2, reduce the lower bounds, thus allowing DynaMIX to operate without errors under tighter deadlines and less memory capacity. Therefore, we can improve system robustness by using the various bit-widths the system supports.

VII. CONCLUSION AND FUTURE WORK

We have proposed a resource optimization framework, called DynaMIX, that allows multiple DNN-based real-time apps to meet the same deadline of apps (dictated by camera’s frame rate), which is critically important for CPS apps. DynaMIX creates the app performance profiles offline, and reconfigures the mixed-precision DNN models at runtime while reflecting dynamically varying resource requirements of the running apps according to their status so as to share

the platform resources. Our implementation in a simulated environment has demonstrated its feasibility and effectiveness.

We assume enough computation resources for the apps so that their models can fit in memory. How to relax this assumption is part of our future work. We would also like to apply DynaMIX in mixed criticality system where multiple apps of different criticality levels can run. Although its current version focuses on CNN-based apps, DynaMIX can be extended to support various DNN-based apps including language models, facilitating the execution of many real-time apps on AVs.

REFERENCES

- [1] Tesla, "https://www.tesla.com/autopilot."
- [2] Toyota, "https://global.toyota/en/newsroom/corporate/35063150.html."
- [3] M. Yang, S. Wang, J. Bakita, T. Vu, F. D. Smith, J. H. Anderson, and J.-M. Frahm, "Re-thinking cnn frameworks for time-sensitive autonomous-driving applications: Addressing an industrial challenge," in *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2019, pp. 305–317.
- [4] S.-C. Lin, Y. Zhang, C.-H. Hsu, M. Skach, M. E. Haque, L. Tang, and J. Mars, "The architectural implications of autonomous driving: Constraints and acceleration," in *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*, 2018, pp. 751–766.
- [5] S. Han, H. Shen, M. Philipose, S. Agarwal, A. Wolman, and A. Krishnamurthy, "Mcdnn: An approximation-based execution framework for deep stream processing under resource constraints," in *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, 2016, pp. 123–136.
- [6] L. N. Huynh, Y. Lee, and R. K. Balan, "Deepmon: Mobile gpu-based deep learning framework for continuous vision applications," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, 2017, pp. 82–95.
- [7] B. Fang, X. Zeng, and M. Zhang, "Nestdnn: Resource-aware multi-tenant on-device deep learning for continuous mobile vision," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 115–127.
- [8] M. Liu, X. Ding, and W. Du, "Continuous, real-time object detection on mobile devices without offloading," in *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2020, pp. 976–986.
- [9] S. Heo, S. Cho, Y. Kim, and H. Kim, "Real-time object detection system with multi-path neural networks," in *2020 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2020, pp. 174–187.
- [10] W. Kang, S. Chung, J. Y. Kim, Y. Lee, K. Lee, J. Lee, K. G. Shin, and H. S. Chwa, "Dnn-sam: Split-and-merge dnn execution for real-time object detection," in *2022 IEEE 28th Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2022, pp. 160–172.
- [11] E. Baek, D. Kwon, and J. Kim, "A multi-neural network acceleration architecture," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 940–953.
- [12] W. Jiang, Z. Song, J. Zhan, Z. He, X. Wen, and K. Jiang, "Optimized co-scheduling of mixed-precision neural network accelerator for real-time multitasking applications," *Journal of Systems Architecture*, vol. 110, p. 101775, 2020.
- [13] J. S. Jeong, J. Lee, D. Kim, C. Jeon, C. Jeong, Y. Lee, and B.-G. Chun, "Band: coordinated multi-dnn inference on heterogeneous mobile processors," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, 2022, pp. 235–247.
- [14] M. Han, H. Zhang, R. Chen, and H. Chen, "Microsecond-scale preemption for concurrent gpu-accelerated dnn inferences," in *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, 2022, pp. 539–558.
- [15] R. Krishnamoorthi, "Quantizing deep convolutional networks for efficient inference: A whitepaper," *arXiv preprint arXiv:1806.08342*, 2018.
- [16] H. Park and C. Kim, "Design of adaptive web and lazy loading components for web application development," *Journal of the Korea Institute of Information and Communication Engineering*, vol. 23, no. 5, pp. 516–522, 2019.
- [17] J. W. Arendt, P. P. Giangarra, R. K. Manikundalam, D. R. Padgett, and J. M. Phelan, "System and method for lazy loading of shared libraries," Jan. 13 1998, uS Patent 5,708,811.
- [18] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [19] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: towards real-time object detection with region proposal networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1137–1149, 2016.
- [20] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [21] K. Wang, Z. Liu, Y. Lin, J. Lin, and S. Han, "Hq: Hardware-aware automated quantization with mixed precision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8612–8620.
- [22] Y. Cai, Z. Yao, Z. Dong, A. Gholami, M. W. Mahoney, and K. Keutzer, "Zeroq: A novel zero shot quantization framework," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13 169–13 178.
- [23] Z. Dong, Z. Yao, A. Gholami, M. W. Mahoney, and K. Keutzer, "Hawq: Hessian aware quantization of neural networks with mixed-precision," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 293–302.
- [24] D. Lee, M. Cho, S. Lee, J. Song, and C. Choi, "A novel sensitivity metric for mixed-precision quantization with synthetic data generation," in *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2021, pp. 1294–1298.
- [25] W. Zhe, J. Lin, V. Chandrasekhar, and B. Girod, "Optimizing the bit allocation for compression of weights and activations of deep neural networks," in *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019, pp. 3826–3830.
- [26] NVIDIA. [Online]. Available: <https://developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth/>
- [27] P. K. Raha, T. Knopp, S. Ahmad, A. Ansari, F.-H. Ho, T. To, V. Nalluri, M. Sarmah, and R. Patwari, "8.2 a versatile 7nm adaptive compute acceleration platform processor," in *2020 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, 2020, pp. 146–148.
- [28] C. Tang, H. Zhai, K. Ouyang, Z. Wang, Y. Zhu, and W. Zhu, "Arbitrary bit-width network: A joint layer-wise quantization and adaptive inference approach," *arXiv preprint arXiv:2204.09992*, 2022.
- [29] A. Bulat and G. Tzimiropoulos, "Bit-mixer: Mixed-precision networks with runtime bit-width selection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5188–5197.
- [30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [32] T. Contributors, "https://pytorch.org/docs/stable/quantization.html," 2019.
- [33] F. Pedregos, "https://github.com/pythonprofilers/memory_profile."
- [34] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6023–6032.
- [35] X. Ding, X. Zhang, N. Ma, J. Han, G. Ding, and J. Sun, "Repvgg: Making vgg-style convnets great again," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13 733–13 742.
- [36] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009.
- [37] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [38] H. Touvron, A. Vedaldi, M. Douze, and H. Jégou, "Fixing the train-test resolution discrepancy," *Advances in neural information processing systems*, vol. 32, 2019.
- [39] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in

context,” in *European conference on computer vision*. Springer, 2014, pp. 740–755.

- [40] W. Kang, S. Chung, J. Y. Kim, Y. Lee, K. Lee, J. Lee, K. G. Shin, and H. S. Chwa, “Dnn-sam: Split-and-merge dnn execution for real-time object detection.”