

# MultiTASC++: A CONTINUOUSLY ADAPTIVE SCHEDULER FOR EDGE-BASED MULTI-DEVICE CASCADE INFERENCE

Sokratis Nikolaidis<sup>1</sup>, Stylianos I. Venieris<sup>2</sup>, Iakovos S. Venieris<sup>1</sup>

<sup>1</sup>National Technical University of Athens, Athens, Greece, <sup>2</sup>Samsung AI Center, Cambridge, UK

NOTE: Corresponding author: Sokratis Nikolaidis, sokratisnikolaidis@mail.ntua.gr

**Abstract** – Cascade systems, consisting of a lightweight model processing all samples and a heavier, high-accuracy model refining challenging samples, have become a widely-adopted distributed inference approach to achieving high accuracy and maintaining a low computational burden for mobile and IoT devices. As intelligent indoor environments, like smart homes, continue to expand, a new scenario emerges, the multi-device cascade. In this setting, multiple diverse devices simultaneously utilize a shared heavy model hosted on a server, often situated within or close to the consumer environment. This work introduces MultiTASC++, a continuously adaptive multi-tenancy-aware scheduler that dynamically controls the forwarding decision functions of devices to optimize system throughput while maintaining high accuracy and low latency. Through extensive experimentation in diverse device environments and with varying server-side models, we demonstrate the scheduler's efficacy in consistently maintaining a targeted satisfaction rate while providing the highest available accuracy across different device tiers and workloads of up to 100 devices. This demonstrates its scalability and efficiency in addressing the unique challenges of collaborative DNN inference in dynamic and diverse IoT environments.

**Keywords** – Cascade, DNN inference, edge computing, multi-tenancy, scheduler

## 1. INTRODUCTION

In the last few years, there have been notable advancements in the realm of on-device execution for Deep Learning (DL) inference tasks [1]. Concurrently, the proliferation of indoor intelligent environments [2], encompassing smart homes and offices, presents an opportunity for DL to facilitate novel applications across a diverse range of smart devices like IoT cameras and AI speakers. However, due to their compact form-factor and energy-consumption constraints, the majority of these devices fall within the low-end of the computational spectrum. Unlike high-end smartphones equipped with robust processors and accelerators such as GPUs and NPUs [3], low-end devices lack the capability to deploy state-of-the-art Deep Neural Networks (DNNs). Consequently, they resort to lightweight models, albeit with lower accuracy.

Considering the drawbacks associated with offloading data to the cloud for inference, such as increased bandwidth usage, latency, and privacy concerns, an alternative strategy has been gaining ground. This new strategy involves placing the server within or in close proximity to the consumer environment, of-

ten in the form of a dedicated AI hub designed to assist nearby devices [2]. Within this framework, cascade architectures have emerged as a notable deployment approach [4, 5, 6, 7, 8, 9]. These architectures capitalize on the inherent variability in sample difficulty, opting to process only the more challenging cases with a robust server-based model, while delegating the processing of simpler samples, which typically constitute the majority of the data stream, to on-device execution using a lightweight model. A substantial body of research has delved into cascade architectures, with a primary focus on refining the forwarding decision criteria and optimizing the selection of model pairs to enhance overall efficacy.

Despite the advancements made, the prevailing focus of existing work has been confined to scenarios where a *single* device utilizes the server at any given moment. This assumption is no longer applicable in the context of emerging intelligent environments, where multiple devices concurrently undertake DL inference tasks with the support of a shared AI hub [10]. This scenario introduces the novel setting of *multi-device cascade*, wherein multiple devices utilize the same model on a shared edge-based

server. A system operating in this mode must exhibit scalability with regard to the number of devices, effectively balancing rapid response times and high accuracy across the devices. Conventional approaches, which treat each model cascade independently, would either lead to brute-forcing inference requests through the server's resources, resulting in system overload, or force all devices to resort to on-device execution, nullifying any accuracy advantages. Consequently, there is a pressing need for innovative methods explicitly tailored to address the challenges posed by a multi-device cascade.

In this setting, the current state-of-the-art work, MultiTASC [11], introduced a multi-tenancy-aware scheduling approach for multi-device cascades. MultiTASC proposed a scheme where forwarding decision functions can be dynamically reconfigured at runtime, providing the adaptability missing from prior cascade architectures. The decision to modify the forwarding functions would be based on the monitoring of the server-hosted model's batch size, regarded as a metric for the server's load. As such, the scheduler would tune the decision functions when the running batch size deviated from a predefined optimal value that was calculated during initialization.

Despite its improvements over conventional cascade systems, MultiTASC resulted in an overly relaxed policy before a certain number of devices was reached and an excessively strict one after the influx of requests became considerably high. Furthermore, when run across multiple independent runs, the observed system behavior demonstrated noticeable variance, indicating that a more fine-grained monitoring approach was needed to ensure a robust deployment.

In this work, we propose MultiTASC++, a continuously adapting multi-tenancy-aware scheduler specifically designed to address the challenges inherent in deploying multi-device cascade architecture in high-demand, AI-enabled indoor spaces. Building upon the strengths of MultiTASC, we retain the dynamically reconfigurable forwarding functions across client devices, aiming to control the server's inference request rate at runtime. Departing from MultiTASC, we introduce a new method for tuning the decision functions that allows for more fine-grained and continuous adaptation. The key contributions of this paper are the following:

- A system model of the multi-device cascade architecture. By expanding the cascade architecture to accommodate multiple devices, our parametrization exposes the tunable parameters and enables system designers to systematically investigate its trade-offs.
- A new multi-tenancy-aware scheduler optimized for the multi-device cascade architecture. With its enhanced approach of reconfiguring the forwarding decision functions, we consider each device's latency requirements independently leading to more effective, device-tailored adaptation. We further introduce the continuous, rather than in discrete steps, tuning of the decision functions, resulting in finer-grained adaptability. Lastly, we introduce server model switching, where the server-side model can be dynamically swapped for another with a different latency-accuracy trade-off. In this manner, we add a new design dimension in the multi-device cascade architecture, further increasing its adaptability.

In the following section, we present the current state of AI-focused edge computing and discuss related work. In Section 3, we describe the system architecture in the multi-device cascade setting as well as our formulation of the target problem. Section 4 presents the novel MultiTASC++ scheduler and its internal design, followed by the experimental evaluation in Section 5. We conclude with Section 6, where we summarize the proposed approach and outline possible avenues for future research.

## 2. RELATED WORK

### 2.1 On-device DNN inference

In recent years, there has been an explosion of Artificial Intelligence (AI) applications and services thanks to significant advancements in DL. These applications span a wide range, from personal assistants and recommendation systems to autonomous vehicles and healthcare diagnostics. Furthermore, the widespread adoption of mobile computing and the Internet of Things (IoT) has led to billions of interconnected mobile and IoT devices, collectively generating an immense volume of data at the network edge [12]. This has created the need to push the execution of DNN applications at the edge of the network leading to a substantial surge in the deployment of DL models on resource-constrained




Fig. 1 – Example of an AI-driven smart office

devices [1]. Although on-device training of DL models remains a challenging endeavor due to the limited computational resources and memory constraints of such devices, significant progress has been made in on-device inference. Several techniques have been proposed to enable efficient on-device inference, including:

**Lightweight model design:** One approach to enable on-device inference is the design of lightweight DL models. Models like MobileNetV2 [13], EfficientNet-Lite [14] and NasNet-Mobile [15] have been specifically crafted to achieve high accuracy while minimizing computational and memory requirements.

**Model quantization:** Quantization techniques [16] reduce the precision of model weights and activations, effectively decreasing memory and computational demands without substantial loss in accuracy.

**Model pruning:** Pruning methods, such as channel pruning [17], aim to reduce the size of DNNs by removing unimportant neurons, thereby reducing computational overhead.

**Knowledge distillation:** Knowledge distillation [18] involves training a smaller, more efficient model to mimic the predictions of a larger, complex model. This allows for the transfer of knowledge from larger models to smaller ones.

**Optimized scheduling:** Scheduling and runtime optimization techniques [19] help allocate computational resources efficiently, ensuring that DL inference tasks run smoothly on constrained devices.

Despite these advancements, modern intelligent environments like smart homes and offices are often equipped with small-form-factor, resource-constrained devices (e.g., smart cameras, AI speak-

ers). These devices lack the processing power to support high-accuracy, computationally-intensive models, driving the need for distributed collaborative inference approaches.

## 2.2 Distributed collaborative inference

Distributed collaborative inference systems leverage a central server to assist mobile and embedded devices in performing DL inference tasks. Notably, the server can be strategically placed at the network edge, close to the devices, to minimize latency and optimize real-time processing. Two primary schemes have emerged in this domain: offloading and cascading.

### A. Offloading

Offloading techniques aim to distribute the computational load between the device and the server. In this scheme, the DNN is divided into two parts, with the initial part executed on the device and the latter part on the server. One standout contribution in this domain is the Neurosurgeon [20] framework, which focuses on the selection of a singular split point for offloading Convolutional Neural Networks (CNNs) from devices to servers, with the objective of optimizing either latency or energy consumption. Later work such as [21, 22] explore the trade-off between latency and accuracy that is introduced when taking the offloading decision. [23, 24] try to address the offloading dilemma in a progressive manner, requiring some training before deployment. Offloading approaches have managed to alleviate some of the computational burdens on the device while maintaining the accuracy of a complex model.

## B. Cascades

Cascade schemes involve a sequence of DNNs with progressively increasing complexity and accuracy. After processing input data through a model, a forwarding decision function determines whether to continue with the current result or proceed to the next, more complex model. [4] is one of the first notable contributions on this approach. It introduces a static technique, calculating an optimized threshold before runtime, as well as a dynamic technique, aiming to optimize the threshold at runtime. The forwarding function in both techniques is defined by analyzing the differences between the softmax results. From our experiments, [4] leads to polarized execution in the majority of cases, with the data being processed either solely locally by the light model or on the server side by the heavy model, *i.e.* without a balanced split between device and server. [5] proposed a trainable forwarding criterion based on a neural head attached to the light model’s feature extractor. This approach yields great results when tested on a certain pair of CNNs but requires training before deployment for every combination of networks. [6] investigates the use of multiple DNN models going beyond two element cascades, coupled with the evaluation of various decision metrics. Additionally, solutions have been proposed to deploy cascades under tight energy constraints [7, 8]. Most of the aforementioned work focus on the image classification task and don’t consider the problem of dynamic adaptation at runtime. The task of video classification has also been studied [25, 26] and lately, the first attempts to accommodate Large Language Models have emerged [27].

### 2.3 Multi-device cascades

Previous research on cascade architectures is predominantly centered around isolated scenarios, where a single device enjoys exclusive access to a dedicated server. Nevertheless, this assumption no longer aligns with the reality of AI-driven indoor environments like the one in Fig. 1. The pervasive integration of AI technologies has given rise to an expanding array of AI-powered devices, resulting in a pressing demand for simultaneous support from a shared server.

In these complex and interconnected scenarios, a more nuanced examination is indispensable. Hasty or simplistic deployment strategies risk overload-

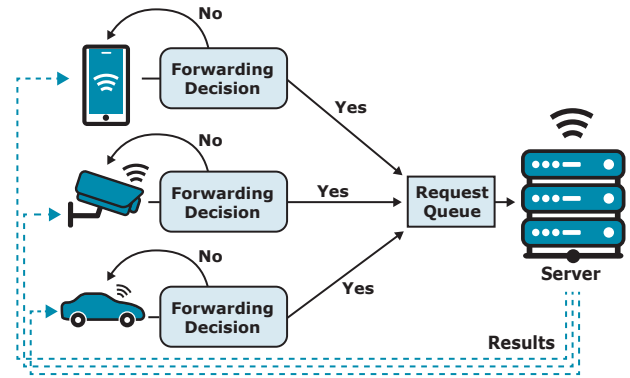


Fig. 2 – System architecture of a multi-device cascade [11]

ing the server and causing protracted response times. However, relying solely on local execution can severely compromise the overall accuracy of the system, as it misses out on the collaborative potential of cascaded processing. Thus, it becomes evident that as we navigate the intricacies of AI-driven indoor environments, thoughtful strategies are paramount to harnessing the full potential of these technologies, optimizing both efficiency and accuracy.

This paper addresses the unexplored setting of multi-device cascades, where multiple devices operate simultaneously, sharing an edge server assisting in the execution of DL inference tasks. This paper provides a principled approach to tackling the challenges of resource allocation and model selection in this complex scenario, facilitating the straightforward and adaptable deployment of such an architecture.

## 3. MULTI-DEVICE INFERENCE CASCADE

In Fig. 2, we present the comprehensive system architecture of a multi-device cascade, specifically designed for executing DL inference tasks on IoT devices in a collaborative setting. Within this architecture, all IoT devices are engaged in performing a common task, such as object detection, albeit they may host different DL models tailored to their computational capabilities and requirements. The main components of this system architecture include the following.

**IoT devices:** These devices are the primary endpoints where the DL inference tasks are executed. Each IoT device is equipped with its own DL model designed to process incoming data efficiently. These devices generate predictions independently based on their respective models.



**Device-hosted models:** DL models on the devices are trained on the same task, operating independently in all other aspects. Each model can be of different architecture, complexity and efficiency, depending on the computational resources and demands of the respective device. Thus, it is of paramount importance that the scheduler maintains high performance in scenarios of device heterogeneity.

**Forwarding decision function:** The output predictions generated by each IoT device are first given to a forwarding decision function. This decision function assesses the confidence of the DL model's output on each device. If the model is sufficiently confident in its prediction, the result remains unchanged, and no further action is taken. However, if there is uncertainty or low confidence in the prediction, the sample is earmarked for further analysis.

**Server:** Samples that require additional scrutiny are forwarded to a centralized server for in-depth processing. The server hosts a more accurate and computationally-intensive DL model capable of refining the predictions made by the IoT devices.

**Request queue:** The forwarded samples from all IoT devices are temporarily stored in a request queue at the server. This queue serves as a staging area where samples awaiting processing are collected. The request queue ensures efficient and organized data flow from the IoT devices to the server.

**Server-hosted model:** The server-side model processes the samples drawn from the request queue. This model is shared among all connected IoT devices, allowing for collaborative refinement of predictions by leveraging the advanced "knowledge" of the network.

**Result distribution:** Finally, the results produced by the server-side model are distributed back to their corresponding IoT devices as soon as they become available. This seamless distribution of refined predictions ensures that each IoT device benefits from the improved accuracy achieved by the server.

### 3.1 Single-device cascade

Let us consider a single IoT device running a classification-based DL inference task. Let  $x \in \mathcal{X}$  be the input and  $y \in \{1, \dots, K\}$  the classification label produced by the model, where  $K$  is the number of classes. The confidence of the given DL model on its output  $y$  can be calculated by using

the decision function  $d(\cdot)$ . The output of  $d(\cdot)$  is binary; for  $d(\cdot) = 0$  we conclude that the model is confident and therefore the result is acceptable, whereas for  $d(\cdot) = 1$  the sample  $x$  should be forwarded to the server for further processing. Denoting the classification function of the light model by  $f_l : \mathcal{X} \rightarrow [0, 1]^K$  that yields the softmax output vector of the model whose maximum value is the predicted class, and the classification function of the heavy model by  $f_h : \mathcal{X} \rightarrow [0, 1]^K$ , we formally define a collaborative cascade system as:

$$casc_{f_l, f_h, d}(x) = \begin{cases} f_l(x) & \text{if } d(f_l(x)) = 0 \\ f_h(x) & \text{if } d(f_l(x)) = 1 \end{cases}$$

### 3.2 Multi-device cascade

To capture multi-device cascade architectures (Fig. 2), we expand the existing single-device cascade system representation as follows. Let  $\mathcal{D}$  be the set of devices assisted by the same server. Then, the multi-device cascade system is defined as:

$$casc_{f_l^i, f_h, d^i}(x^i) = \begin{cases} f_l^i(x^i) & \text{if } d^i(f_l^i(x^i)) = 0 \\ f_h(x^i) & \text{if } d^i(f_l^i(x^i)) = 1 \end{cases} \quad \forall i \in \{1, \dots, |\mathcal{D}|\}$$

where  $x^i \in \mathcal{X}^i$  is a sample processed by the  $i$ -th device,  $f_l^i$  the classification function of the DL model deployed on the  $i$ -th device,  $f_h$  the shared heavy model on the server, and  $d^i(f_l^i(x^i))$  the forwarding decision function of the  $i$ -th device. Note that the shared heavy model  $f_h$ , is the only variable not depending on the devices.

### 3.3 Congestion problem

In the case of the single-device cascade, the server's computational resources are exclusively assisting a single device, resulting in minimal response time and preservation of accuracy. However, in the area of IoT spaces it is increasingly important to leverage a server's capabilities across multiple devices simultaneously, hence implementing a multi-device cascade. This approach paves the way towards amortizing the server's cost and maximizing its utility. Nonetheless, depending on the specific conditions, if the arrival rate of incoming requests exceeds the server's processing throughput capacity, the server becomes overwhelmed, leading to extended waiting times for requests in the queue.

For a given number of devices denoted by  $|\mathcal{D}|$ , Eq. (1) expresses the arrival rate of requests to the

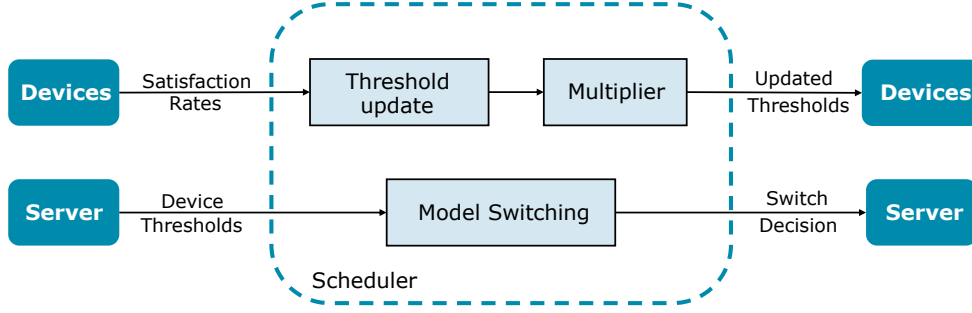


Fig. 3 – Architecture of the MultiTASC++ scheduler

server, *i.e.* the rate at which results are deemed as unsatisfactory by the decision functions on the devices and are therefore forwarded to the server.

$$AR_{\text{server}} = \sum_{i=1}^{|\mathcal{D}|} \frac{p_{\text{casc}}^i}{t_{\text{inf}}^i} \quad (1)$$

where  $t_{\text{inf}}^i$  is the average inference latency of a sample on the  $i$ -th device and  $p_{\text{casc}}^i$  is the probability of a sample giving  $d^i(f_i^i(x^i)) = 1$ .

Given the attainable throughput  $T_{\text{server}}$  of the server, we distinguish three different states:

- $AR_{\text{server}} < T_{\text{server}}$ : The processing rate of the server exceeds the arrival rate, leading to the server being underutilized. Forwarding a greater quantity of challenging samples to the server could enhance accuracy.
- $AR_{\text{server}} = T_{\text{server}}$ : A state of equilibrium is reached, where requests are promptly processed upon arrival, preventing accumulation, and ensuring full utilization of the server's processing power.
- $AR_{\text{server}} > T_{\text{server}}$ : The rate of incoming requests surpasses the server's processing capacity. If this condition persists, it will result in a substantial accumulation of requests in the queue, leading to excessive latency.

The  $p_{\text{casc}}^i$  of the forwarding decision function is not static since it depends on the processing order of the samples. Therefore, this architecture involves stochastic components at runtime and could greatly benefit by dynamically adapting its state depending on the current conditions. Since  $t_{\text{inf}}^i$  and  $T_{\text{server}}$  are fixed based on the device and server-side processors, we opt to manipulate  $p_{\text{casc}}^i$  by changing the parameters of  $d^i(f_i^i(x^i))$  in order to introduce adaptability to the system.

### 3.4 Problem optimization

We frame the aforementioned setting as a multi-objective optimization problem, seeking to maximize accuracy and throughput subject to a latency Service-Level Objective (SLO). The following section describes our proposed scheduler designed to address it.

## 4. PROPOSED SCHEDULER

To address the challenges related to request accumulation and efficient server resource utilization, we introduce MultiTASC++, a multi-tenancy-aware scheduler that dynamically adapts the arrival rate of samples from the assisted devices. Based on our previous work, MultiTASC, we kept the component of reconfigurable forwarding decision functions to control the arrival rate and introduced four new techniques that completely change the way the thresholds are updated: *i)* SLO satisfaction rate updates, *ii)* continuous threshold reconfiguration, *iii)* threshold scaling, and *iv)* server model switching. MultiTASC++'s internal architecture is presented in Fig. 3 where all new techniques are visible.

### 4.1 Reconfigurable forwarding decision function

Significant research endeavors have been dedicated to assessing the prediction confidence of DNNs, resulting in the development of various approaches [6, 31]. In this work, we employ the Best-versus-Second-Best (BvSB metric [32], which quantifies prediction confidence by computing the difference between the top two values in the softmax output of the model (referred to as  $P_1$  and  $P_2$ )) as shown in Eq. (2). These values correspond to the highest and

**Table 1** – Evaluated DNN models

Model	Location Device		Clock Rate	Accuracy	Latency	FLOPs	#Params
MobileNetV2 [13]	Low-end	Sony Xperia C5	1.69 GHz	71.85%	31 ms	0.6 B	3.5 M
EfficientNetLite0 [14]	Mid-tier	Samsung A71	2.20 GHz	75.02%	43 ms	0.8 B	4.7 M
EfficientNetB0 [14]	High-end	Samsung S20 FE	2.73 GHz	77.04%	33 ms	0.8 B	5.3 M
MobileViT-x-small [28]	High-end	Google Pixel 7	2.85 GHz	74.64%	57 ms	1.1 B	2.3 M
InceptionV3 [29]	Server	Tesla T4 GPU	585 MHz	78.29%	15 ms	11.4 B	23.8 M
EfficientNetB3 [14]	Server	Tesla T4 GPU	585 MHz	81.49%	25 ms	3.7 B	12.2 M
DeiT-Base-Distilled [30]	Server	Tesla T4 GPU	585 MHz	83.41%	14 ms	7.7 B	86.0 M

\* See Table 1 in [1] for the detailed resource characteristics of the target mobile phones.

second-highest classes that the classifier predicted.

$$\text{BvSB} \Big|_{f(x)} = P_1 - P_2 \quad (2)$$

Other metrics, such as top-1 softmax or entropy can be implemented in the system with minimal modifications, potentially leading to different latency-accuracy trade-offs.

In contrast to the predominant approach employed by most existing cascade systems, which establish fixed decision thresholds during design and maintain them upon deployment, our work introduces an alternative approach. We adopt a *dynamic reconfiguration* scheme for the decision function to cater to the adaptability requirements of our target system. The decision function  $d^i(\cdot)$  is defined as shown in Eq. (3), with adjustments facilitated by a dynamic scheduler that fine-tunes its parameters at runtime.

$$d^i(f_l^i(x)) = \begin{cases} 0 & \text{if } \text{BvSB} \Big|_{f_l^i(x)} \geq c_{i,t} \\ 1 & \text{if } \text{BvSB} \Big|_{f_l^i(x)} < c_{i,t} \end{cases} \quad (3)$$

where  $c_{i,t}$  is the decision threshold of device  $i$  at time  $t$ . The per-device decision thresholds are exposed to our server-residing scheduler, which adapts them at runtime.

#### 4.2 SLO satisfaction rate updates

Given a latency target, we introduce the SLO satisfaction rate metric as the percentage of samples successfully processed within the designated latency constraint. Latency, in this context, is measured from the initiation of inference on the device until the final result is obtained, either by the device-hosted model or the server-side model in the cases where the sample is forwarded to the

server. The SLO satisfaction rate relies significantly on the timely processing of samples forwarded to the server, a factor influenced by the volume of samples that the server must handle.

The SLO satisfaction rate is the main metric we use to measure the smooth operation of the system since a high satisfaction rate can be a good indicator of a responsive queue without congestion. By having a metric that helps us understand the state of the system during runtime, we can tune the influx of samples with high precision by adjusting the reconfigurable forwarding decision functions of each device individually.

Each device calculates the average SLO satisfaction rate during its inference in time windows of  $T$  seconds. At the end of every time window, the satisfaction rate for that window is forwarded to the server where the appropriate reconfiguration of the decision function is calculated. By constantly informing the server about the satisfaction rate, we achieve a beneficial trade-off of accuracy in order to maintain the desired SLO satisfaction rate.

#### 4.3 Continuous threshold reconfiguration

Every  $T$  seconds, the device sends an SLO satisfaction rate update to the scheduler. Assuming there is a target SLO satisfaction rate value depending on the user’s needs, if the update value is below that optimal value, the server chooses to reduce the number of samples coming from that device. Otherwise, if the update value is greater than the designated optimal value, the server chooses to increase the influx of samples from that device to achieve a higher accuracy.

MultiTASC++ leverages the information available

through the SLO satisfaction rate updates to view the forwarding decision thresholds as continuous variables. This allows for incredible precision in its ability to identify the optimal threshold for the situation and pair of models, as well as rapid adaptation when needed. The completely revamped update rule is presented in Eq. (4) and is in stark contrast to MultiTASC's [11] update rule where the scheduler had to guess the optimal influx of samples and then use a very slow and imprecise step based approach to converge to it.

$$\Delta thresh = -a \cdot (SR_{target} - SR_{update}) \quad (4)$$

where  $\Delta thresh$  is the amount by which the threshold will be adjusted,  $SR_{target}$  is the target SLO satisfaction rate for that device,  $SR_{update}$  is the SLO satisfaction rate sent by the device and  $a$  is a scaling factor.

#### 4.4 Threshold scaling

The update rule presented in Eq. (4) demonstrates robust performance in cases where the optimal threshold is either lower or close to the initial threshold of the device. However, its responsiveness diminishes when facing scenarios characterized by substantial underutilization of server resources, necessitating a rapid threshold increase. To accommodate such cases, we introduced a scaling component to the threshold update. Alg. 1 presents the proposed approach. Following the threshold update from Eq. (4), the updated threshold is subsequently either scaled by a multiplier  $m$  (line 4) if the threshold was increased with the update (*i.e.*  $SR_{update} > SR_{target}$ ), or stays the same otherwise. In the first case,  $m$  is then updated by the rule shown in line 9 while in the second case it is reset to 1. Since for a large number of devices the multiplier is not required, we further incorporate a penalty term (line 9) that considers the number of devices currently active in the system, denoted by  $n$  (line 5).

#### 4.5 Server model switching

To further improve the adaptability of MultiTASC++, we introduce the server model switching feature. Depending on the dynamic conditions, the scheduler can decide to switch to a different server-hosted model whose computational cost-accuracy trade-off better suits the current state. For instance, if a heavy model is being used and the optimal SLO

---

#### Algorithm 1 Multiplier Implementation

---

##### Input:

- 1:  $SR_{target}$  ▷ The targeted satisfaction rate
- 2:  $SR_{update}$  ▷ The device's satisfaction rate
- 3:  $thresh_{updated}$  ▷ The updated threshold
- 4:  $m$  ▷ Previous multiplier
- 5:  $n$  ▷ Number of active devices

##### Output:

- 6:  $thresh_{final}$  ▷ Threshold sent to device
  - 7: **if**  $SR_{target} < SR_{update}$  **then**
  - 8:      $thresh_{final} = m \cdot thresh_{updated}$
  - 9:      $m = m \cdot (1 + \frac{0.1}{n})$
  - 10: **else**
  - 11:      $thresh_{final} = thresh_{updated}$
  - 12:      $m = 1$
  - 13: **end if**
- 

satisfaction rate can't be maintained without significantly dropping the accuracy, MultiTASC++ opts to switch to a faster model to allow for a greater influx of samples. On the opposite, if a faster model is being used but the system is low load, resulting in underutilization of the server, MultiTASC++ chooses to switch to a heavier model to achieve a higher accuracy without sacrificing the SLO satisfaction.

The scheduler makes a decision by examining the current thresholds of the devices. To switch from a heavy model to a faster one, every device's threshold in a single tier must be below a certain value. On the other hand, to switch from the fast model to the heavier one, every device's threshold must be above a certain value depending on its tier.

Given the set of all device thresholds  $\mathcal{C}$  consisting of elements  $c_i^k \in [0, 1]$ , where  $k$  is the device tier (*i.e.*  $k \in \mathcal{K} = \{low, mid, high\}$ ) and  $i \in \mathcal{D}^k$  corresponds to the  $i$ -th device in this tier, we can define the model switch decision as follows: for  $S(\mathcal{C}) = -1$  a switch to a faster model is needed, for  $S(\mathcal{C}) = +1$  a switch to a heavier model is needed and for  $S(\mathcal{C}) = 0$  the current model is the optimal. Formally:

$$S(\mathcal{C}) = \begin{cases} -1 & \text{if } \exists k \in \mathcal{K} : c_i^k < c_{lower}, \forall i \in \mathcal{D}^k \\ +1 & \text{if } c_i^k > c_{upper}, \forall k \in \mathcal{K}, \forall i \in \mathcal{D}^k \\ 0, & \text{otherwise} \end{cases}$$

Upper and lower limits  $c_{upper}^k$  and  $c_{lower}$  are set after a thorough examination of cascade results on a training set.

## 5. EVALUATION

### 5.1 Experimental setup

To evaluate the performance of MultiTASC++, we built a prototype in Python 3.9 utilizing TensorFlow 2.9.1. The edge server component hosts an NVIDIA Tesla T4 GPU, Intel(R) Xeon(R) 2.30GHz CPU and 12GB of RAM, allowing it to run heavy, state-of-the-art models with high accuracy within the evaluated latency targets. On the device components' side, we used mobile devices spanning three different tiers, namely high, mid, and low-end, using Samsung S20 FE, Samsung A71 and Sony Xperia C5 Ultra, respectively. We also used the Google Pixel 7 as a high-end device for the purpose of evaluating transformer models. For on-device execution, we used TensorFlow Lite and targeted the CPU of the respective mobile device, as it is still the most widely-used approach [3, 33].

Different models were chosen for each tier, with the aim of having the highest accuracy possible while taking into account the device's computational resources. We measured the average inference latency of each model on the respective device across 200 runs with a batch size of 1. We followed the same process to measure the average server inference latency across different batch sizes and used this data to conduct simulation-based experiments. The experiments targeted a variety of scenarios and mainly focused on the SLO latency target and accuracy metrics. Communication between the devices and the edge server component was established using the AMPQ protocol, following the widely-used practice for communication between IoT devices. The protocol was implemented through the AMPQS-*form* library which allows for thread-safe execution.

To fully take advantage of the server's computational resources and boost throughput, it is important to use batching, *i.e.* processing multiple samples at the same time. To avoid the latency that would arise from waiting for the request queue to reach a specific batch size, we employ dynamic batching [34]. With dynamic batching, we use the maximum batch size that is feasible with the current request queue length. Available batch sizes are  $\mathcal{B}=\{1, 2, 4, 8, 16, 32, 64\}$ . Due to diminishing returns, in some cases we use a lower maximum batch size, *e.g.* with EfficientNetB3 a batch size of 16 provides a higher throughput and lower latency than a batch size of 32 and above.

**Models and datasets:** In our experiments, we target the task of image classification with 1k classes. Concretely, we used the 50k-images validation set of the ImageNet dataset [35]. Table 1 presents the evaluated models, as well as their accuracy and inference latency on the different execution points. We obtained the CNN models from TensorFlow Hub while the transformer models from Hugging Face. All models are pretrained on ImageNet's training set. We deploy MobileNetV2, EfficientNetLite0 and EfficientNetB0 to the low, mid and high-end client devices respectively. On the server side, we use InceptionV3 and EfficientNetB3 to explore the differences between a lower-accuracy, higher-throughput model and a higher-accuracy, lower-throughput model, respectively.

We also evaluated our scheduler using transformer models both as device and server-hosted models. We deployed MobileVit-x-small on Google Pixel 7 and DeiT-Base-Distilled on the server. As the transformer architecture has only recently started gaining traction, available models are not as efficient as their CNN counterparts on most mobile devices. As such, we utilize a computationally-powerful flagship phone, like Pixel 7, to achieve acceptable latency. Nevertheless, as transformer models become more efficient [36], they are expected to soon become broadly deployed across mobile devices.

**Evaluation settings:** To assess our system across deployment setups, we focused on two distinct cases: *i)* a homogeneous scenario, which comprises devices of equal processing capabilities that host the same local model, and *ii)* a heterogeneous scenario, which comprises devices of diverse processing capabilities, with each device hosting a model tailored to its respective tier. We also conduct experiments as part of the homogeneous scenario to evaluate the transformer models, the server model switching technique, as well as a scenario that emulates a realistic setting of intermittent device participation, *i.e.* where devices drop in and out during inference.

For general evaluation in the homogeneous scenario, all devices are of the same tier, using Sony Xperia C5 Ultra with MobileNetV2. We selected this tier, as it represents the configuration with the minimum latency, thereby imposing the greatest






Fig. 4 – SLO satisfaction rate for InceptionV3 - MobileNetV2




Fig. 5 – Accuracy for InceptionV3 - MobileNetV2




Fig. 6 – Throughput for InceptionV3 - MobileNetV2

challenge on the scheduler. Additionally, it accentuates the substantial disparity in accuracy between the device-hosted model and the full cascade accuracy, effectively showcasing the cascade architecture’s potential. In the heterogeneous scenario, all three tiers of devices were deployed in equal percentage. Finally, to assess the scheduler’s performance when using transformer models, we used a separate tier of Pixel 7 devices running MobileViT- x-small.

In all scenarios but one, the dataset of each device consisted of 5000 randomly selected samples from the last 40000 images of ImageNet’s validation set. The exception consisted of a distinct scenario where

the dataset consisted of only 1000 samples. We followed this approach in order to more clearly highlight a limitation of MultiTASC [11] where the SLO satisfaction rate was below its required value due to slow convergence of the threshold reconfiguration. As shown later in Fig. 10, MultiTASC++ alleviates effectively this issue.

Three different seeds were used to sample the data and the average values of each metric, alongside their minimum and maximum, are reported. The metrics used for the evaluation are: the system throughput capturing the system’s processing rate in samples/s, the average accuracy across devices, the latency SLO satisfaction rate for 100, 150 and




Fig. 7 – SLO satisfaction rate for EfficientNetB3 - MobileNetV2




Fig. 8 – Accuracy for EfficientNetB3 - MobileNetV2




Fig. 9 – Throughput for EfficientNetB3 - MobileNetV2

200 ms SLOs, and the scalability in terms of number of devices.

**Baselines:** As a baseline, we use Static, a scheduler with statically selected thresholds that remain fixed during runtime. To tune the static threshold, we use the first 10000 images of ImageNet’s validation set as our calibration set and evaluate all cascade model pairs in terms of accuracy and forwarding probability. As such, we tune the threshold so that approximately 30% of samples are forwarded to the heavy model, providing a balanced accuracy-latency trade-off. In cases where that threshold yielded an accuracy loss of more than 1 pp compared to the highest achievable cascade accuracy,

we used the lowest threshold that satisfied the 1 pp limit. This baseline is equivalent to a set of state-of-the-art cascades [5, 6, 9]. We also compare with MultiTASC [11], which constitutes the current state-of-the-art scheduling method for multi-device cascades.

### 5.2 Evaluation of performance

In this section, we assess the performance of our scheduler compared to Static and MultiTASC across all scenarios. We set the target SLO value to 95, meaning that we aim for 95% of samples to finish inference within the latency target, independent of whether they are forwarded to the server or stay




Fig. 10 – EfficientNetB3 - MobileNetV2 with 1000 samples

on the device. We do not aim for 100% so that the system has some leeway to trade SLO satisfaction rate percentage for accuracy. Here, we highlight a significant improvement introduced by MultiTASC++. Our new scheduler allows us to set and consistently maintain a target satisfaction rate regardless of the prevailing conditions, addressing a limitation present in MultiTASC.

We also set the time window  $T$  to 1.5 s and the scaling variable  $a$  of the continuous threshold to 0.005. Compared to MultiTASC, the amount of variables to initialize and the computational effort required to determine their values have notably decreased. This streamlines the scheduler’s deployment process, while delivering substantially improved and consistent results. Furthermore, the new continuous threshold update scheme allows for the SLO targets to be chosen independently for each device, contrasting MultiTASC where all devices had to agree on the same latency target during the initialization of the scheduler.

#### A. Homogeneous scenario

Fig. 4 shows the SLO satisfaction rate as the number of devices increases with InceptionV3 hosted on the server. The figure includes the MultiTASC++, Static and MultiTASC approaches. Notably, MultiTASC++ manages to consistently maintain the satisfaction rate close to or above 95%. In contrast, Static experiences rapid degradation, with the majority of the results of the forwarded samples not returning within the latency target for 25-40 devices and above. On the other hand, MultiTASC exhibits a dip in the range of approximately 5 to 40 devices, followed by an overcorrection that achieves a 100% satisfaction rate resulting in a needless degradation

in accuracy. This dip occurs due to MultiTASC opting to use batch size as a metric for congestion, which proves to be suboptimal. Pairing that with a static step update rule, the scheduler was neither able to accurately predict the state of the system nor adapt with the speed that is required in borderline system states. This is mitigated in MultiTASC++, where the scheduler is more accurately aware of the state of the system through the SLO updates from the devices (Section 4.2), while also being able to adapt more quickly through the techniques of continuous threshold reconfiguration (Section 4.3) and threshold scaling (Section 4.4).

Fig. 5 presents the full cascade’s accuracy in comparison to the baselines as the number of devices increases. MultiTASC++ effectively strikes a better balance between accuracy and the maintenance of the desired 95% SLO satisfaction rate. In comparison to Static, MultiTASC++ achieves a higher accuracy for a smaller number of devices where the server is being underutilized, while for larger amounts of devices, it chooses to trade off accuracy in order to sustain the satisfaction rate. Importantly, even though the accuracy is lower compared to the static approach, it is still substantially higher than the device-hosted model’s accuracy, justifying the use of a cascade architecture. Compared to MultiTASC, our scheduler achieves higher accuracy across all cases, with the exception of the instances where MultiTASC’s satisfaction rate performance (Fig. 4) dips below the desired 95%.

Fig. 6 shows the system throughput difference between MultiTASC++ and Static as the number of devices increases. This figure shows the need for dynamic runtime adaptation, since, while Static stagnates at 1000 samples per second, MultiTASC++




Fig. 11 – SLO satisfaction rate for InceptionV3 - different devices

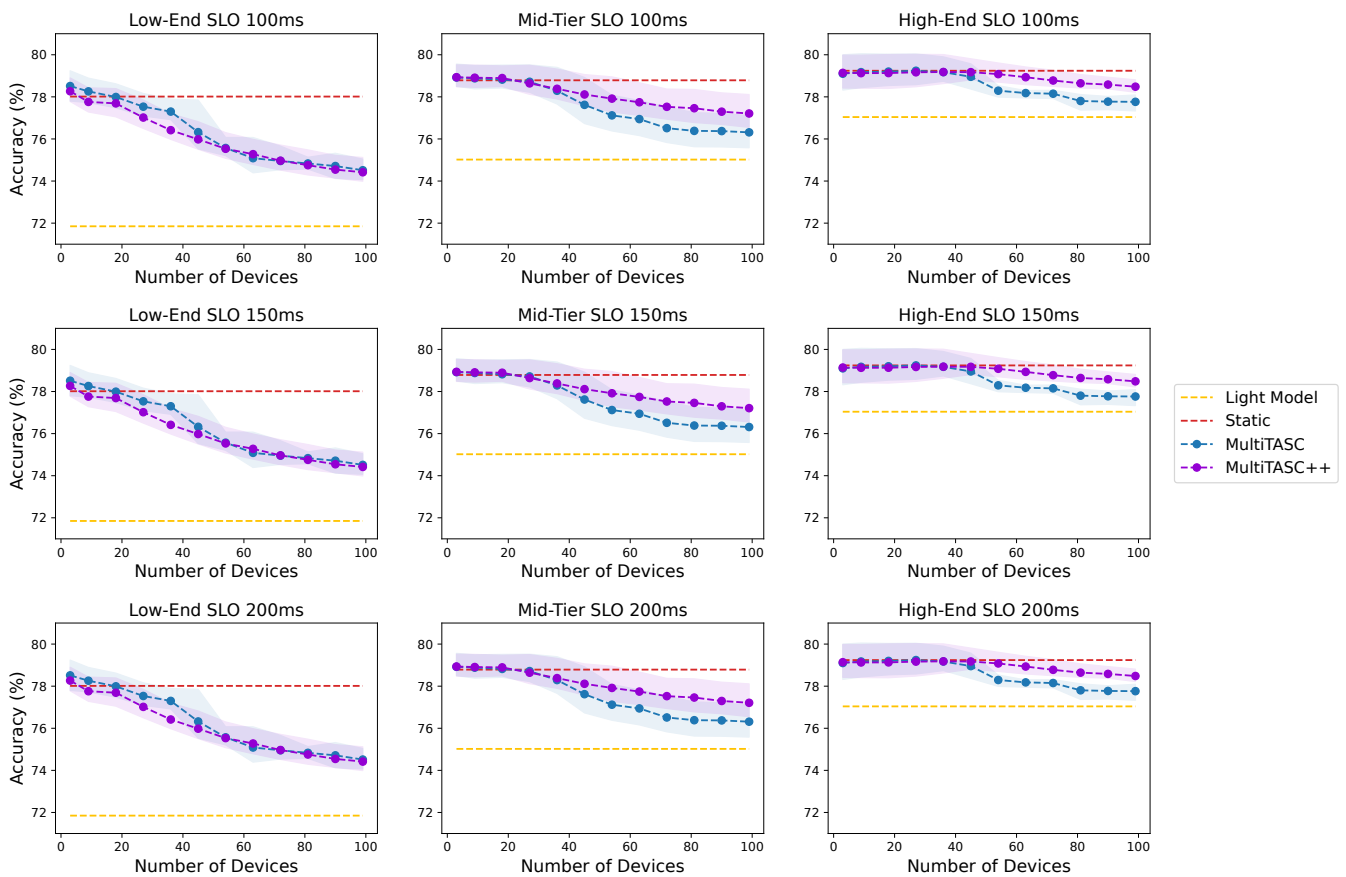


Fig. 12 – Accuracy for InceptionV3 - different devices

manages to keep the linear increase of system throughput as the number of devices rises for all SLO targets.

Fig. 7 and Fig. 8 depict the achieved SLO satisfaction rate and accuracy, when EfficientNetB3 is deployed on the server. A similar dip to the one in Fig. 4 can be observed for MultiTASC, between approximately 5 and 20 devices. The dips in both figures reach values as low as 80%, which represents unacceptable delays in 15pp more samples than the target. We can also observe a more significant differ-

ence in accuracy, especially when targeting a 100 ms SLO, due to the lower attainable throughput of the heavier EfficientNetB3. Notably, even with a substantial number of devices, the collaborative cascade architecture, when paired with MultiTASC++, significantly improves upon the accuracy of the on-device model while preserving consistent responsiveness. In contrast to MultiTASC, besides optimizing the trade-off between satisfaction rate and accuracy, our approach also minimizes the variance between different seed runs, a crucial element in ensuring a robust high-quality service. Fig. 9 shows




Fig. 13 – SLO satisfaction rate for EfficientNetB3 - different devices

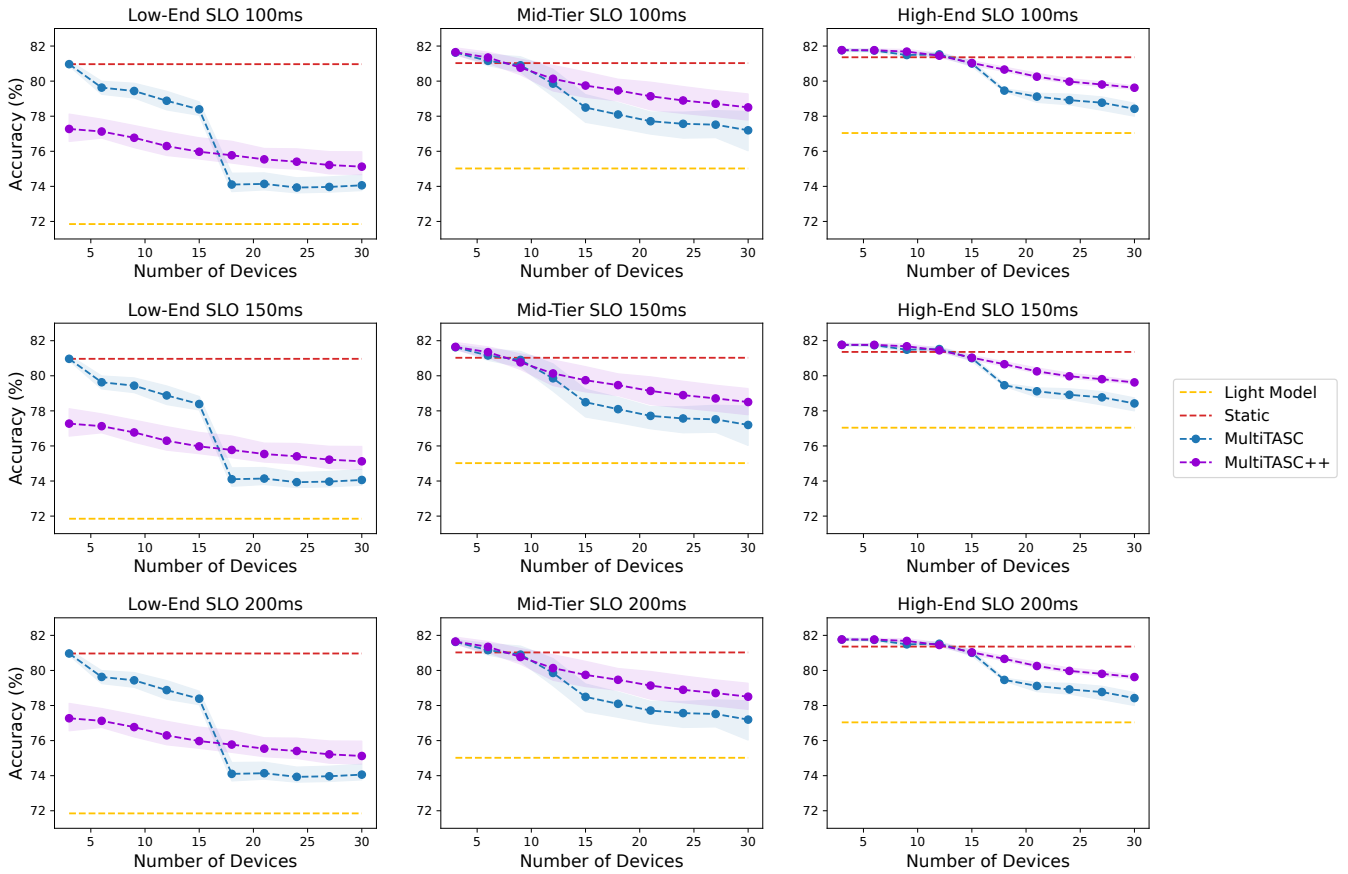


Fig. 14 – Accuracy for EfficientNetB3 - different devices




Fig. 15 – SLO satisfaction rate for DeiT-Base-Distilled - MobileViT-x-small






Fig. 16 – Accuracy for DeiT-Base-Distilled - MobileViT-x-small

similar results to Fig. 6, where Static converges to a system throughput of around 300 samples per second while MultiTASC++ manages to maintain a linear increase of system throughput as the number of devices increases.

Fig. 10 shows the SLO satisfaction rate and accuracy trends with an increasing number of devices, utilizing a reduced dataset of 1000 samples, as opposed to the previous 5000 and with a lenient 150 ms SLO. A noticeable distinction becomes apparent when comparing the evaluated approaches. MultiTASC converges slowly to a threshold that satisfies the SLO, evident by the results between 10 to 20 devices, where the satisfaction rate reaches a low point of 75%. In contrast, MultiTASC++ consistently delivers nearly identical results to those observed in the prior experiment.

*B. Heterogeneous scenario*

Fig. 11 and Fig. 12 present a comprehensive comparison of SLO satisfaction rate and accuracy across MultiTASC++, MultiTASC, and Static in a heterogeneous device environment. In this setting, InceptionV3 serves as the server-side model. We report the performance metrics separately for each device tier.

Similar to the homogeneous scenario, our observations underscore the limitations of the static approach, which experiences a critical failure beyond a certain number of devices. Conversely, both MultiTASC++ and MultiTASC effectively maintain high satisfaction rates. Notably, our proposed scheduler exhibits superior consistency and efficiency in trading accuracy for satisfaction rate, as it robustly maintains satisfaction rates precisely at the targeted level. MultiTASC, while providing operational viability compared to Static for more devices, it also introduces significant variance

that can undermine the quality of service. Additionally, it does not effectively utilize the given satisfaction rate allowance, resulting in lower accuracy levels than potentially achievable.

Furthermore, MultiTASC++ mitigates the dip experienced by MultiTASC within the range of approximately 5 to 40 devices. In particular, when it comes to mid and high-end tier devices, MultiTASC++ achieves significant accuracy gains over MultiTASC, while accuracy remains relatively consistent for low-end devices between both schedulers.

Fig. 13 and Fig. 14 present the same comparison with EfficientNetB3 as the server-side model. The results and conclusions obtained are similar to those from the previous experiment, where InceptionV3 was the server model.

However, it is worth emphasizing that the variance and fluctuations observed in MultiTASC’s satisfaction rate are even more pronounced in this scenario. In terms of accuracy, MultiTASC achieves a higher accuracy with low-end devices, particularly in the range of 0 to 15 devices, but subsequently lags behind, mirroring the patterns observed in mid and high-end devices. This phenomenon can be attributed to the low-end devices’ lower latency, which significantly contributes to congestion, necessitating faster threshold adjustments. It is also worthy to note that MultiTASC++ achieves a higher accuracy for light devices when it comes to larger numbers of devices.

*C. Evaluation on transformers*

Fig. 15 and Fig. 16 offer an examination of SLO satisfaction rates and accuracy in response to an in-




Fig. 17 – Model switching with InceptionV3 initialization

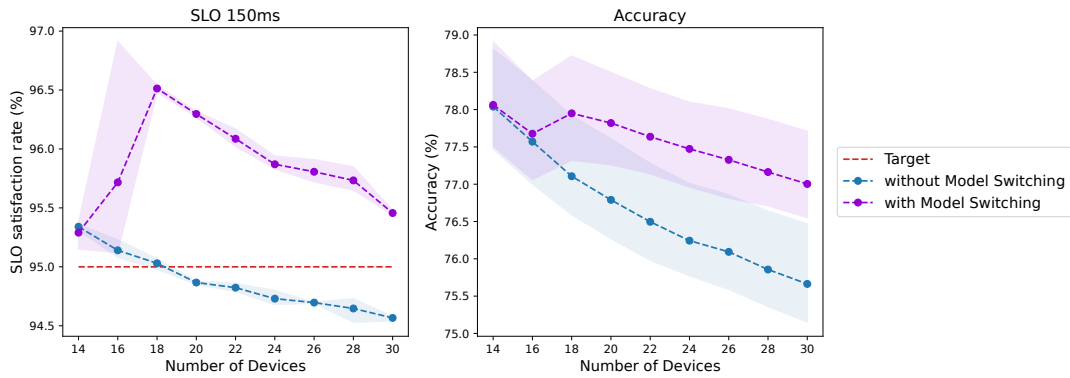


Fig. 18 – Model switching with EfficientNetB3 initialization




Fig. 19 – MultiTASC++’s behavior under intermittent device participation with dynamic threshold

creasing number of devices using MultiTASC++ and Static. In this scenario, the device model is the mobile-grade MobileViT-x-small transformer, while the server model is DeiT-Base-Distilled.

The outcomes closely resemble those observed in previous scenarios, showcasing the ability of our scheduler to generalize to transformer architectures, even though its design was not tailored to

them. Any apparent disparities in the results compared to the CNN-based experiments can be attributed to differences in inference latencies and model accuracies.

#### D. Evaluation of model switching

Fig. 17 presents a comparative analysis of the SLO satisfaction rate and accuracy with the server model switching feature enabled and disabled. The scheduler is initialized with the InceptionV3 model, targeting a 150 ms latency. Up to 12 devices, during runtime MultiTASC++ decides to switch to a heavier model, namely EfficientNetB3. This achieves a substantially higher accuracy while keeping the satisfaction rate above the 95% target. For 14 devices and above, the switch is no longer efficient and InceptionV3 is used across all samples. We should note here that the model switching feature was not used in previous experiments so that our update rule could be fairly evaluated against Multi-TASC without other techniques affecting the performance.

Fig. 18 shows a similar evaluation for the scheduler, with the server-side model initialized to EfficientNetB3 and targeting a 150 ms latency. Once again, the satisfaction rate consistently exceeds the 95% target, accompanied by an observable accuracy enhancement when the model switches to InceptionV3. In scenarios featuring 14 and 16 devices, model switching does not occur, leading to outcomes similar to instances without model switching.

#### E. Intermittent device participation

In this experiment, we emulate a realistic setting where 20 devices run simultaneously, each bearing a 50% probability of going offline. We target low-tier devices with EfficientNetB3 as the server-hosted model. The point at which a device goes offline follows a normal distribution with a mean  $\mu = N/2$  and a standard deviation  $\sigma = N/5$  where  $N$  is the total number of samples. The duration for which a device remains offline adheres to an alpha distribution with a shape parameter  $\alpha = 60$  seconds.

Fig. 19 shows an overview of the dynamics in this experiment. It illustrates the fluctuation in the number of active devices over time, the average threshold maintained across devices, the running SLO sa-

tisfaction rate, and the average running accuracy for the currently active devices. We note that the percentage of active devices, the average threshold, and the running satisfaction rate are represented on the left y-axis scale, while the running accuracy is on the right y-axis scale. Several key observations emerge from this visualization.

First, we notice an inverse correlation between the threshold and the number of active devices. Initially, the threshold rapidly decreases to accommodate the need for maintaining the satisfaction rate. Subsequently, as the number of active devices diminishes, the threshold is increased.


Furthermore, a direct correlation is observed between the increase in threshold and the gain in running accuracy. This emphasizes the trade-off mechanism facilitated by the scheduler to optimize system performance.

Notably, the running SLO satisfaction rate consistently maintains a level exceeding approximately 95% throughout the duration of the experiment, underlining MultiTASC++'s effectiveness in ensuring the fulfillment of the specified service-level objectives.

Fig. 20 shows the results of another experiment with different initializations when it comes to devices, leading to a different plot line representing the number of active devices compared to the previous experiment. In this scenario, however, a static threshold is employed. The results from this experiment unveil several insights.

The utilization of a static threshold, set at 0.35, engenders relatively stable running accuracy, as a completely static running accuracy is not realistic. However, it is accompanied by a notable variability in the running satisfaction rate, consistently falling well below the 95% target. This observation underlines the critical role played by a dynamically adaptive threshold in ensuring that SLOs are met.

Both the running accuracy and the running satisfaction rate eventually converge at approximately 250 seconds when the devices complete their inference tasks. Meanwhile, the threshold and the number of active devices continue to vary. This divergence arises due to extensive congestion within the request queue, resulting in a delay of approximately 30 seconds between the devices completing their in-



**Fig. 20** – MultiTASC++'s behavior under intermittent device participation with static threshold

ference and the server's return of all requested results. Since approximately 50% of the devices do not go offline, we expect the number of active devices to drop in half at around 155 seconds where their inference finishes. While this can be observed in Fig. 19 where the system remains responsive due to MultiTASC++, in the case of Fig. 20 this does not happen. This is due to the congestion in the request queue and the devices staying connected, waiting for the results to return from the server.

Overall, this examined scenario further highlights the benefit of dynamic threshold adaptation during runtime, particularly in settings characterized by evolving conditions and varying system demands.

## 6. CONCLUSION

In this paper, we presented MultiTASC++, a dynamic scheduler designed to address the challenges of congestion in collaborative DNN inference involving a multitude of IoT devices in indoor intelligent spaces. By introducing the concept of the multi-tenant cascade, we have achieved continuous dynamic adaptation of threshold values, optimizing the trade-off between accuracy and service-level objectives during runtime.

Our experimental evaluation, spanning diverse device environments and server-side models, has demonstrated the efficacy of our dynamic scheduler. Moreover, these experiments comparing Mul-

tiTASC++ to our old implementation, led us to the conclusion that the SLO satisfaction rate updates are critical for precisely adhering to the targets set by each device. Equally essential is the adoption of continuous threshold values, paired with a method designed to counteract slow updates, for the scheduler to achieve its full potential. The harmonic cooperation between all of these components led to a scheduler characterized by its consistent ability to maintain the satisfaction rate that is targeted, as well as allowing for system throughput to scale, while simultaneously enhancing performance across various device tiers and workloads, surpassing the abilities of its predecessor, MultiTASC.

Moreover, the implementation of model switching has proven effective, allowing for the scheduler to adapt to scenarios where a different server-hosted model leads to greater efficiency.

Lastly, by emulating a realistic deployment scenario with intermittent partial participation of devices, we demonstrated our dynamic scheduler's versatility, which can adapt to dynamic device availability and tune the execution configuration in a timely manner. This adaptability is vital for optimizing system performance and ensuring that the service-level objectives are met.

While image classification is a well-explored and commonplace task, it would be of interest for future research to shift attention towards other significant

tasks, such as speech recognition. Furthermore, the proposed system could be extended by investigating the new challenges of generative inference tasks, such as image and text generation, and designing novel methods in order to enable their deployment in multi-device cascade setups.

## REFERENCES

- [1] Stylianos I. Venieris, Ioannis Panopoulos, and Iakovos S. Venieris. "OODIn: An Optimised On-Device Inference Framework for Heterogeneous Mobile Devices". In: *SMARTCOMP*. 2021.
- [2] Stefanos Laskaridis, Stylianos I Venieris, Alexandros Kouris, Rui Li, and Nicholas D Lane. "The Future of Consumer Edge-AI Computing". In: *arXiv* (2022).
- [3] Mario Almeida, Stefanos Laskaridis, Abhinav Mehrotra, Lukasz Dudziak, Ilias Leontiadis, and Nicholas D. Lane. "Smart at What Cost? Characterising Mobile Deep Neural Networks in the Wild". In: *IMC*. 2021.
- [4] Eunhyeok Park, Dongyoung Kim, Soobeom Kim, Yong-Deok Kim, Gunhee Kim, Sungroh Yoon, and Sungjoo Yoo. "Big/little deep neural network for ultra low power inference". In: *CODES+ISSS*. 2015.
- [5] Min Li, Yu Li, Ye Tian, Li Jiang, and Qiang Xu. "AppealNet: An Efficient and Highly-Accurate Edge/Cloud Collaborative Architecture for DNN Inference". In: *DAC*. 2021.
- [6] Xin Wang, Yujia Luo, Daniel Crankshaw, Alexey Tumanov, Fisher Yu, and Joseph E Gonzalez. "IDK Cascades: Fast Deep Learning by Learning not to Overthink". In: *UAI*. 2018.
- [7] Seyed Iman Mirzadeh and Hassan Ghasemzadeh. "Optimal Policy for Deployment of Machine Learning Models on Energy-Bounded Systems". In: *IJCAI*. 2020.
- [8] Dimitrios Stamoulis, Ting-Wu Rudy Chin, Anand Krishnan Prakash, Haocheng Fang, Sribhuvan Sajja, Mitchell Bogner, and Diana Marculescu. "Designing Adaptive Neural Networks for Energy-Constrained Image Classification". In: *ICCAD*. 2018.
- [9] Alexandros Kouris, Stylianos I Venieris, and Christos-Savvas Bouganis. "CascadeCNN: Pushing the Performance Limits of Quantisation in Convolutional Neural Networks". In: *FPL*. 2018.
- [10] Tetsuro Nakamura, Shogo Saito, Kei Fujimoto, Masashi Kaneko, and Akinori Shiraga. "Spatial- and Time-Division Multiplexing in CNN Accelerator". In: *Parallel Computing* (2022).
- [11] Sokratis Nikolaidis, Stylianos I. Venieris, and Iakovos S. Venieris. "MultiTASC: A Multi-Tenancy-Aware Scheduler for Cascaded DNN Inference at the Consumer Edge". In: *2023 IEEE Symposium on Computers and Communications (ISCC)*. 2023.
- [12] Zhi Zhou, Xu Chen, En Li, Liekang Zeng, Ke Luo, and Junshan Zhang. "Edge Intelligence: Paving the Last Mile of Artificial Intelligence With Edge Computing". In: *Proceedings of the IEEE* (2019).
- [13] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. "MobileNetV2: Inverted Residuals and Linear Bottlenecks". In: *CVPR*. 2018.
- [14] Mingxing Tan and Quoc V. Le. "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks". In: *International Conference on Machine Learning (ICML)*.
- [15] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. "Learning Transferable Architectures for Scalable Image Recognition". In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
- [16] Song Han, Huizi Mao, and William J Dally. "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding". In: *ICLR*. 2015.
- [17] Yihui He, Xiangyu Zhang, and Jian Sun. "Channel Pruning for Accelerating Very Deep Neural Networks". In: *ICCV*. 2017.
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. "Distilling the Knowledge in a Neural Network". In: *NeurIPS*. 2014.



- [19] Bart Cox, Robert Birke, and Lydia Y. Chen. “Memory-aware and Context-aware Multi-DNN Inference on the Edge”. In: *Pervasive and Mobile Computing* (2022).
- [20] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. “Neurosurgeon: Collaborative Intelligence Between the Cloud and Mobile Edge”. In: *ASPLOS*. 2017.
- [21] Hongshan Li, Chenghao Hu, Jingyan Jiang, Zhi Wang, Yonggang Wen, and Wenwu Zhu. “JALAD: Joint Accuracy-And Latency-Aware Deep Structure Decoupling for Edge-Cloud Execution”. In: *ICPADS*. 2018.
- [22] Mario Almeida, Stefanos Laskaridis, Stylianos I Venieris, Ilias Leontiadis, and Nicholas D Lane. “DynO: Dynamic Onloading of Deep Neural Networks from Cloud to Device”. In: *TECS* (2022).
- [23] Stefanos Laskaridis, Stylianos I. Venieris, Mario Almeida, Ilias Leontiadis, and Nicholas D. Lane. “SPINN: Synergistic Progressive Inference of Neural Networks over Device and Cloud”. In: *MobiCom*. 2020.
- [24] Jin Huang, Colin Samplawski, Deepak Ganesan, Benjamin Marlin, and Heesung Kwon. “CLIO: Enabling Automatic Compilation of Deep Learning Pipelines across IoT and Cloud”. In: *MobiCom*. 2020.
- [25] Daniel Kang, John Emmons, Firas Abuzaid, Peter Bailis, and Matei Zaharia. “NoScope: Optimizing Neural Network Queries over Video at Scale”. In: *Proceedings of the VLDB Endowment* (2017).
- [26] Haichen Shen, Seungyeop Han, Matthai Philipose, and Arvind Krishnamurthy. “Fast Video Classification via Adaptive Cascading of Deep Models”. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
- [27] Yiding Wang, Kai Chen, Haisheng Tan, and Kun Guo. “Tabi: An Efficient Multi-Level Inference System for Large Language Models”. In: *EuroSys*. 2023.
- [28] Sachin Mehta and Mohammad Rastegari. “MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer”. In: *The Tenth International Conference on Learning Representations, ICLR*. 2022.
- [29] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. “Rethinking the Inception Architecture for Computer Vision”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- [30] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. “Training data-efficient image transformers & distillation through attention”. In: *International Conference on Machine Learning (ICML)*. 2021.
- [31] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. “On Calibration of Modern Neural Networks”. In: *ICML*. 2017.
- [32] Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. “Multi-Class Active Learning for Image Classification”. In: *CVPR*. 2009.
- [33] Ioannis Panopoulos, Sokratis Nikolaidis, Stylianos I. Venieris, and Iakovos S. Venieris. “Exploring the Performance and Efficiency of Transformer Models for NLP on Mobile Devices”. In: *IEEE Symposium on Computers and Communications (ISCC)*. 2023.
- [34] Ahsan Ali, Riccardo Pinciroli, Feng Yan, and Evgenia Smirni. “BATCH: Machine Learning Inference Serving on Serverless Platforms with Adaptive Batching”. In: *SC*. 2020.
- [35] Li Fei-Fei, Jia Deng, and Kai Li. “ImageNet: Constructing a Large-Scale Image Database”. In: *Journal of Vision* 9.8 (2009), pp. 1037–1037.
- [36] Junting Pan, Adrian Bulat, Fuwen Tan, Xiatian Zhu, Lukasz Dudziak, Hongsheng Li, Georgios Tzimiropoulos, and Brais Martinez. “EdgeViTs: Competing Light-Weight CNNs on Mobile Devices with Vision Transformers”. In: *European Conference on Computer Vision (ECCV)*. 2022.

## AUTHORS



**Sokratis Nikolaidis** is currently a Ph.D. candidate in the School of Electrical and Computer Engineering, National Technical University of Athens, since 2021, and an assistant researcher at Intelligent Communications and Broadband Networks Laboratory.

He received an M.Eng. degree in applied mathematical and physical sciences in 2021 from National Technical University of Athens. He has published two research papers for the international conference ISCC 2023, receiving the Best Student Full Paper award. His current research interests include distributed ML/DL, computer vision, natural language processing and generative AI.



**Stylianos I. Venieris** is currently a senior research scientist at Samsung AI, Cambridge, U.K., where he leads the Distributed AI group. He received a Ph.D. degree in reconfigurable hardware and deep learning in 2018 and

the M.Eng. degree (Hons.) in electrical and electronic engineering in 2014 from Imperial College London, London, U.K. He has published over 20 research papers in peer-reviewed journals and international conferences. His latest research interests include custom hardware accelerators for deep neural networks, highly flexible inference systems for mobile and embedded platforms, and continual federated learning.



**Iakovos S. Venieris** has been a professor at the School of Electrical and Computer Engineering, National Technical University of Athens (NTUA) since 1994 and is the director of Intelligent Communications and Broadband Networks (ICBNet) Laboratory.

His research interests are in the fields of distributed systems, security and privacy, software and service engineering, agent technology, multimedia, mobile communications and machine learning, internetworking, signalling, resource scheduling and allocation for network management, modeling, performance evaluation, queuing theory. He has over 370 publications in the above areas and has received several national and international awards for academic achievement. He has participated in and has successfully led several national and international projects.