

Sparse-DySta: Sparsity-Aware Dynamic and Static Scheduling for Sparse Multi-DNN Workloads

Hongxiang Fan
hongxiangfan@ieee.org
Samsung AI Center &
University of Cambridge
Cambridge, UK

Stylianos I. Venieris*
s.venieris@samsung.com
Samsung AI Center
Cambridge, UK

Alexandros Kouris
a.kouris@samsung.com
Samsung AI Center
Cambridge, UK

Nicholas D. Lane
ndl32@cam.ac.uk
University of Cambridge &
Flower Labs
Cambridge, UK

ABSTRACT

Running multiple deep neural networks (DNNs) in parallel has become an emerging workload in both edge devices, such as mobile phones where multiple tasks serve a single user for daily activities, and data centers, where various requests are raised from millions of users, as seen with large language models. To reduce the costly computational and memory requirements of these workloads, various efficient sparsification approaches have been introduced, resulting in widespread sparsity across different types of DNN models. In this context, there is an emerging need for scheduling sparse multi-DNN workloads, a problem that is largely unexplored in previous literature. This paper systematically analyses the use-cases of multiple sparse DNNs and investigates the opportunities for optimizations. Based on these findings, we propose *Dysta*, a novel bi-level dynamic and static scheduler that utilizes both static sparsity patterns and dynamic sparsity information for the sparse multi-DNN scheduling. Both static and dynamic components of *Dysta* are jointly designed at the software and hardware levels, respectively, to improve and refine the scheduling approach. To facilitate future progress in the study of this class of workloads, we construct a public benchmark that contains sparse multi-DNN workloads across different deployment scenarios, spanning from mobile phones and AR/VR wearables to data centers. A comprehensive evaluation on the sparse multi-DNN benchmark demonstrates that our proposed approach outperforms the state-of-the-art methods with up to 10% decrease in latency constraint violation rate and nearly 4× reduction in average normalized turnaround time. Our artifacts and code are publicly available at: <https://github.com/SamsungLabs/Sparse-Multi-DNN-Scheduling>.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; • **Theory of computation** → **Scheduling algorithms**; • **Computer systems organization** → **Real-time system architecture**.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MICRO '23, October 28–November 1, 2023, Toronto, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0329-4/23/10...\$15.00
<https://doi.org/10.1145/3613424.3614263>

KEYWORDS

Sparse Multi-DNN Scheduling, Dynamic and Static Approach, Algorithm and Hardware Co-Design

ACM Reference Format:

Hongxiang Fan, Stylianos I. Venieris, Alexandros Kouris, and Nicholas D. Lane. 2023. Sparse-DySta: Sparsity-Aware Dynamic and Static Scheduling for Sparse Multi-DNN Workloads. In *56th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '23)*, October 28–November 1, 2023, Toronto, ON, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3613424.3614263>

1 INTRODUCTION

The unprecedented inference accuracy of DNNs has led to the rapidly increasing co-location of multiple DNN-powered applications [59]. Spanning from mobile [1, 29] and wearable apps [37, 66] to cloud services [6, 56], this trend leads to a significant rise in multi-DNN workloads. Simultaneously, recent advances in model compression have enabled faster and smaller models through various sparsification techniques [44, 61]. With inference workloads being commonly sparse, optimization opportunities have been explored, focusing on single-DNN execution [9]. This emergence of sparsity has also influenced AI hardware, as evidenced by the sparsity support of NVIDIA GPUs via Tensor Core [48] and Google TPU v4 via SparseCore [30]. Nonetheless, despite its ubiquitousness across DNNs, sparsity has remained largely unexplored in multi-DNN workloads. As depicted in Figure 1(a), the increasing number of parallel requests raised by diverse user applications has resulted in a surging demand for efficiently handling multiple sparse DNNs. Therefore, there is an urgent need to study sparse multi-DNN workloads.

One of the key components while executing such multi-tenant workloads is the *scheduler*. Being responsible for deciding which task to dispatch next to the processing engine, extensive evaluation [2, 7, 49] has demonstrated that the quality of scheduling determines to a large degree the attainable performance under DNN multi-tenancy, rendering the design of the scheduler a critical task. Although various multi-DNN scheduling approaches have been proposed [2, 17, 33, 36, 49, 58], they are currently bounded by two main limitations:

- These scheduling approaches are optimized to perform well primarily on a single metric, e.g. either minimizing the average normalized turnaround time (ANTT) [14] or the latency service-level objective (SLO) violations, while excessively degrading the other (Section 2.3.3).
- The sparsity is largely neglected in these previous approaches, leading to suboptimal results. More specifically, these methods

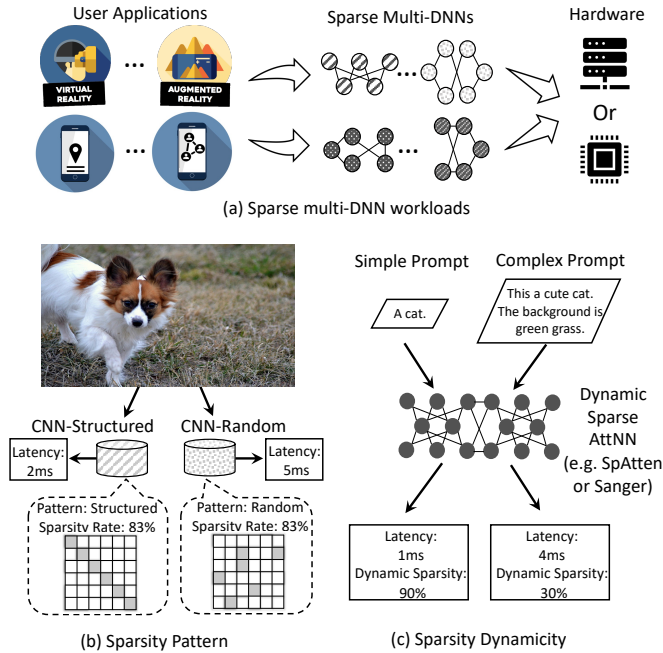


Figure 1: Sparse multi-DNN workloads with different sparsity patterns and runtime dynamics.

do not consider fine-grained details, such as sparsity patterns and dynamicity, that are crucial for further optimizations under sparse multi-DNN workloads.

In this paper, we aim to take the first step in systematically investigating the optimization opportunities in sparse multi-DNN workloads. To achieve this, we propose a public sparse multi-DNN workload benchmark, which includes convolutional neural networks (CNNs) for vision tasks and attention-based neural networks (AttNNs) for natural language processing (NLP) applications. These benchmark models are sparsified using dynamic and static approaches with different patterns in order to emulate real-world scenarios. As shown in Figure 1, we identify two sparsity properties that lead to significant inter-model (across models) and intra-model (across input samples) dynamicity at runtime:

- Sparsity pattern, which refers to the pattern of non-zero masks used while sparsifying weights. Figure 1(a) presents an example of CNNs with the same sparsity rate but different sparsity patterns, yielding varied latencies due to the count of effective operations.
- Sparsity dynamicity, which comprises the input-dependent sparsity in activations that leads to latency variability across input samples. Figure 1(b) illustrates an NLP example, where short and simple prompts with less information may lead to higher dynamic sparsity and lower latency compared to long and complex prompts at runtime. This category of dynamic sparsity is usually generated by either dynamic pruning approaches [20, 43, 52] or certain operations, such as rectified linear units (ReLU) [47], that regularly produce zero values.

These two sparsity properties lead to the dynamic behavior of sparse DNNs during execution, which affects the scheduling design

Table 1: Comparison of existing multi-DNN schedulers.

Design	ANTT Optimized	SLO Viol. Optimized	Dynamic & Static Sparsity	Pattern Aware
PREMA [7]	✓	✗	✗	✗
AI-MT [2]	✓	✗	✗	✗
Layerweaver [49]	✓	✗	✗	✗
Planaria [17]	✗	✓	✗	✗
SDRM ³ [33]	✗	✓	✗	✗
<i>Dysta</i> (Our Work)	✓	✓	✓	✓

while optimizing SLO violation rate and ANTT. We provide further profiling and statistical analysis in Section 2.3.

Based on the identified sources of sparsity dynamicity and the proposed benchmark, this paper proposes *Dysta*, a novel bi-level dynamic and static scheduler that optimizes for both ANTT and violation rate. At the first software-based level, *Dysta* utilizes static information, such as sparsity patterns and average latency, to obtain the initial task priorities before the execution of each task. At the second level, we introduce a lightweight hardware scheduler that adaptively refines the scheduling by monitoring the runtime information, such as the sample-specific dynamic sparsity, for further optimization. Table 1 presents a comparison of existing multi-DNN schedulers. To foster further development of multi-DNN acceleration research, we intend to open-source our code, including benchmarks, evaluation infrastructure and proposed approach, and submit our work for artifact evaluation. Overall, this work makes the following contributions:

- A public benchmark of sparse multi-DNN workloads that contains a wide range of DNNs with various forms of sparsity, including both dynamic and static sparsity and diverse sparsity patterns (Section 3).
- A bi-level dynamic and static scheduling framework, named *Dysta*, which utilizes both static and dynamic sparsity information to optimize the processing of sparse multi-DNN workloads (Section 4).
- A low-cost hardware design of the *Dysta* scheduler and an efficient sparse latency predictor. Several hardware optimizations are proposed to reduce the resource consumption of the scheduler, making its hardware overhead negligible compared with the overall accelerator (Section 5).

To assess the effectiveness of our method, we conducted a comprehensive evaluation on the new benchmark we introduced (Section 6). Our approach achieves significant improvements in both ANTT and violation rate over the state-of-the-art multi-DNN schedulers, demonstrating its advantages and the value of the new benchmark for evaluating schedulers on sparse multi-DNN workloads.

2 BACKGROUND AND MOTIVATION

2.1 Multi-DNN Workloads and Scheduling

The scheduling strategy constitutes a core component of multi-DNN accelerators that directly affects the attainable performance of the system [59]. In the context of multi-DNN workloads with the layer-wise processing manner, the scheduler determines which layer of which model should be processed next and is the main

driver of preemption and resource allocation decisions. Depending on the characteristics of the target application, scheduling can be either static or dynamic. Static scheduling [31, 36, 58] is suitable for fixed-purpose multi-DNN systems, such as robots and autonomous vehicles, where the target set of DNNs and their performance requirements are known *a priori*. In contrast, dynamic scheduling [2, 7, 17, 33, 49] is a more flexible approach, where dynamic schedulers allow new inference tasks to be processed and resources to be repurposed based on the completed and already running DNNs.

Table 1 presents the existing multi-DNN dynamic scheduling algorithms [2, 7, 17, 33, 49]. A key commonality behind them is their strong reliance on the assumption that DNN inference constitutes a predictable workload, *i.e.* it comprises a static computation graph for all inputs. Based on this assumption, execution time estimates, obtained through an offline profiling stage, are then employed to guide runtime scheduling decisions. Nonetheless, the static-workload assumption leads to overly conservative schedules and, in turn, to inefficient utilization of the hardware accelerator. As described in Section 2.3, sparsity constitutes an increasingly ubiquitous form of dynamicity in modern and upcoming DNNs that shifts the status-quo fixed DNN workload towards a dynamic, input-dependent workload, with the sparsity level changing in a per-sample manner. With multi-DNN workloads putting unprecedented pressure on computational demands, there is an emerging need for novel sparsity-aware solutions that extract the maximum performance from the underlying accelerator and enable the scalable processing of multiple DNNs.

2.2 Sparse DNN Accelerators

Stemming from sparsity-inducing training [34], activation functions [47] or pruning methods [3], sparsity is widely observed across different types of models. At the hardware front, a series of works have proposed sparse DNN accelerators [10, 28, 43, 45, 51, 62] that manage to obtain speedup and efficiency gains through zero-skipping mechanisms and efficient sparse-storage schemes. Besides the static weight sparsity that is known *a priori*, advanced sparse accelerators [10, 28, 43, 45, 51] also support dynamically skipping ineffectual computations based on the input-dependent sparsity of activations, providing additional performance improvements over conventional accelerators.

2.3 Motivation

In the context of multi-DNN workloads, previous literature has so far overlooked the various types of sparsity that are observed across models. In this paper, we identify two primary properties of sparsity, namely 1) sparsity dynamicity and 2) sparsity pattern, and study their effects in scheduling sparse multi-DNN workloads.

2.3.1 Sparsity Dynamicity. Sparsity dynamicity is widely observed across different families of DNNs, including both attention-based (AttnNs) and convolutional neural networks (CNNs). The two main sources of dynamicity are *i)* dynamic pruning methods and *ii)* operations that regularly generate zero values.

For AttnNs, the most common form of sparsity dynamicity comes from dynamic pruning approaches. By focusing on weak




Figure 2: Impact of dynamic sparsity on language models.




Figure 3: Sparsity ratios of *ResNet-50* and *VGG-16*.

Table 2: Relative range of network sparsity.

Model	<i>GoogLeNet</i>	<i>VGG-16</i>	<i>InceptionV3</i>	<i>ResNet-50</i>
Relative Range	28.3%	21.8%	23.0%	15.1%

connections between different tokens in the attention matrix, various pruning techniques have been proposed to exploit the attention sparsity for acceleration [19, 20, 43, 52, 62, 70]. To demonstrate the dynamic behavior introduced by the attention sparsity, we profile the per-layer latency of *BERT* [12] on the accelerator introduced by [43]. We iterate the sparse *BERT* over the *SQUAD* [54] dataset. For better visualization, we normalize the distribution by the average latency. As it can be seen in Figure 2, the normalized latency varies from 0.6 to 1.8, resulting in significant dynamic behavior during runtime.

For CNNs, we focus on the ReLU activation function [47] as an example of an operation that regularly generates zeros at runtime and study its dynamic behavior. Although *PREMA* [7] has briefly discussed the impact of ReLU on runtime dynamicity in CNNs, it did not consider out-of-distribution or less informative inputs that may cause large sparsity variances, such as images taken by users in dark or poorly illuminated environments. In this paper, we conduct a more comprehensive analysis by including low-light images from the *ExDark* [42] and *DarkFace* [63] datasets to emulate real scenarios. We profile the activation sparsity of the last six layers in *ResNet-50* and *VGG-16*. As shown in Figure 3, the sparsity ratios of most layers range from 10% to 45%, indicating that a large variance is introduced when considering out-of-distribution or less informative images. To investigate the sparsity of the network instead of individual layers, we calculate the network sparsity by averaging individual layer sparsities across the whole network. As shown in Table 2, the relative range of network sparsity can reach up to 28.4%, depending on the model.




Figure 4: Impact of weight sparsity pattern on the valid MAC operations on *ResNet-50* and *MobileNet*.

2.3.2 Sparsity Pattern. Sparsity pattern refers to the mask type used when sparsifying DNNs. Mainstream sparsity patterns include point-wise random [21], block-wise structured [15, 69] and channel-wise sparsification [23]. The support for different sparsity patterns largely depends on the underlying hardware design. For instance, the Sparse Tensor Core introduced in the NVIDIA Ampere architecture [48] supports the block-wise N:M pattern, while other customized accelerators have better performance for the point-wise random pattern [5, 64]. Although previous work attempts to leverage sparsity information when scheduling multi-DNN workloads, they only rely on the overall sparsity ratio, without considering the sparsity patterns and their relationship with the target hardware.

To investigate the effect of sparsity patterns, we profile the amount of valid MAC operations introduced by point-wise random and channel-wise sparsities, respectively. For both patterns, we keep the same overall sparsity ratio with identical input images for profiling. Figure 4 presents the distribution of normalized MAC operations on *ResNet-50* and *MobileNet* with sparsity ratios of 95% and 80%, respectively. Notably, different sparsity patterns may introduce up to 40% difference in normalized valid MACs, despite the same sparsity ratio.

2.3.3 Optimization Opportunities. Given our analysis, we identify two key optimization opportunities for the scheduler of sparse multi-DNN workloads.

Opportunity One: Dynamicity- & Pattern-Aware — Based on the profiling results of Sections 2.3.1 and 2.3.2, sparsity may lead to significant dynamic behavior at runtime, which affects the optimality of schedulers when targeting sparse multi-DNN workloads. Figure 5 presents an example of the Shortest-Job First (SJF) scheduler, to illustrate the importance of capturing sparsity information for multi-DNN workloads. When exploiting fine-grained information, such as dynamic sparsity ratio or sparsity pattern, the scheduler can make a more informed preemption decision based on accurate latency estimation, avoiding the violation of the second request. Based on this observation, in this work, we aim to utilize sparsity dynamicity and pattern information to improve sparse multi-DNN scheduling.

Opportunity Two: SLO- & ANTT-Optimized — Although previous approaches have attempted to optimize multi-DNN scheduling, we observe that they only perform well on a single metric, e.g. either




Figure 5: SJF schedulings with and without sparsity information in task violation.

Table 3: Benchmark Models.

Scenarios	Applications	Tasks	Models
Data Center	Visual Perception	Object Detection	<i>SSD</i>
		Image Classification	<i>VGG-16, ResNet-50</i>
Mobile Phone	Personal Assistant	Machine Translation	<i>BART, GPT-2</i>
		Question & Answering	<i>BERT</i>
AR/VR	Hand Tracking	Hand Detection	<i>SSD</i>
Wearables		Gesture Recognition	<i>MobileNet</i>

the latency SLO violation rate or ANTT, while sacrificing the performance of the other. For instance, compared with the traditional SJF scheduler, current SOTA approaches can only improve ANTT under the condition of excessively higher SLO violations [7]. The reasons are that they either *i)* are not deadline-aware or *ii)* do not consider the optimization of the violation rate or ANTT while including deadlines. In this paper, we aim to address these drawbacks by co-optimizing both SLO violation rate and ANTT.

3 SPARSE MULTI-DNN BENCHMARK

3.1 Benchmark Models

One of the limitations in previous multi-DNN research is the absence of publicly accessible benchmarks for conducting fair comparisons. To address this issue, we propose a public benchmark for sparse multi-DNN scheduling. Our benchmark is designed to accommodate a diverse collection of sparsified CNNs and AttNNs.

As shown in Table 3, we consider sparse multi-DNN workloads in three setups: data center, mobile phones and AR/VR wearables. Our benchmark comprises six tasks spanning three distinct applications: visual perception, personal assistant, and hand tracking. The tasks include object detection, image classification, machine translation, question & answering, hand tracking and gesture recognition. For vision tasks, we adopt four popular CNN models, namely *SSD* [41], *ResNet-50* [22], *VGG-16* [57], and *MobileNet* [24]. For AttNNs, our benchmark includes three commonly used language models: *BERT* [12], *BART* [39], and *GPT-2* [53]. In terms of datasets, we adopt ImageNet [11], ExDark [42], DarkFace [63], and COCO [40] for training and evaluation on vision tasks, while GLUE [60] and SQUAD [54] are used for language tasks.

3.2 Model Sparsification

In order to study the impact of sparsity patterns, we adopt three different pruning methods for CNNs: random point-wise [21], N:M block-wise [69] and channel-wise pruning [23]. The generated sparsity patterns of each pruning approach are shown in Figure 6. We obtain the pre-trained CNNs from *PyTorch* [50], and apply the target pruning methods using sparsification recipes provided by *SparseML* [35] and *SparseZoo*.^{*} We expose the sparsity rate as a tunable parameter.




Figure 6: Three sparsity patterns applied on CNNs.

For AttNNs, we adopt the dynamic pruning approach introduced in [43] to study the effect of sparsity dynamicity. The adopted method performs pruning dynamically via binary thresholding based on a lightweight prediction of the attention matrix. Following their open-source code,[†] we set the threshold as 0.2 for *BART*, and 0.002 for *BERT* and *GPT2* to maintain the original accuracy of each model.

3.3 Evaluation Methodology

3.3.1 Methodology Overview. Figure 7 shows our evaluation process, comprising hardware simulation and scheduling evaluation phases. During the *Hardware Simulation* phase, we insert our *Python*-based hardware simulator into sparse models using the *Hook* function provided by *PyTorch*. This enables the generation of runtime information such as per-layer latency and sparsity ratio of each model for the target sparse hardware. We obtain the runtime information for each input and model pair by processing the whole dataset with each sparse model. The runtime information is then saved as files for later use. In the *Scheduling Evaluation* phase, we generate multiple task requests by sampling from sparse models summarized in Sections 3.1 and 3.2. The *Scheduler Engine* is then deployed to handle incoming requests according to the designated scheduling algorithm. The runtime information obtained in the hardware simulation phase is used by the *Scheduler Engine* to simulate the execution of the target hardware. After the completion of all jobs, we evaluate different metrics, including both ANTT and violation rate based on the user-specified latency SLO.

3.3.2 Hardware Simulator. To simulate sparse AttNNs, we select *Sanger* [43] as the target hardware accelerator with the support of load-balanced computation for attention sparsity. Their open-source simulator[‡] is adopted in our evaluation. For CNNs, although there are various sparse accelerators, only a few of them open-source their design or simulator. The most promising third-party

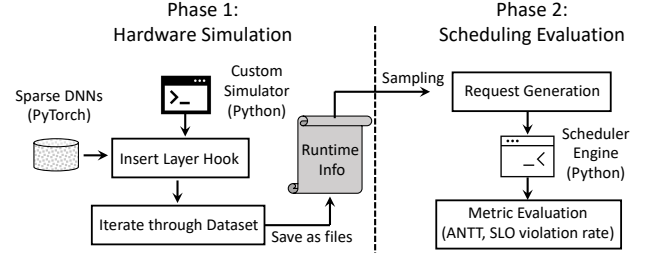


Figure 7: Evaluation methodology.




Figure 8: Overview of *Dysta*'s scheduling approach.

simulators are *SparseLoop* [64] and *STONNE* [46]. However, these frameworks suffer from either prohibitively slow evaluation when actual data are used or limited *PyTorch* support, rendering them unsuitable for meeting our needs. As a result, we develop a custom *Python*-based simulator for sparse CNN evaluation. In this endeavor, we select *Eyeriss-V2* as the target CNN accelerator for two reasons. First, it supports both weight and activation sparsity. Second, several third-party implementations are available for replicating the *Eyeriss-V2* hardware design.[‡] This allowed us to use these references to validate the correctness of our custom simulator.

4 DYNAMIC AND STATIC SCHEDULER

4.1 Framework Overview

As demonstrated in Section 2.3.3, the dynamicity and pattern of sparsity introduce variations in runtime behavior that can significantly impact the optimality of scheduling. To alleviate this, we introduce *Dysta*, a bi-level scheduling method that combines static and dynamic scheduling while leveraging information about both sparsity dynamicity and pattern.

Figure 8 presents an overview of the proposed method. Our bi-level scheduler consists of a static and a dynamic component, implemented at the software and hardware levels, respectively. At the first level, the static component assigns the incoming requests with initial scores to determine the processing order, taking into account the sparsity patterns of the DNN workloads. At this step, the static scheduler also populates a lookup table (LUT) with each request's model information. This per-request information includes *i)* the model's sparsity pattern, *ii)* the average sparsity ratio across all layers, and *iii)* the average latency on the target hardware. The average sparsity and latency information is obtained by profiling

^{*}<https://sparsezoo.neuralmagic.com/>

[†]<https://github.com/hatsu3/Sanger>

[‡]https://github.com/SingularityKChen/dl_accelerator and https://github.com/karthiugumar/CSE240D-Hierarchical_Mesh_NoC-Eyeriss_v2

Algorithm 1: Static Scheduler

```

1 Input:  $Reqst_n = \langle Model_n, Pattn_n, input_n, SLO_n \rangle$ 
2 Output:  $Score_n, Info_n$ 
3 if  $Reqst_n$  arrives then ▷ New request arrives
4    $Info_n \leftarrow LUT(Reqst_n)$  ▷ Obtain model info from LUT
5    $Lat_n \leftarrow PredLat(Model_n, Pattn_n)$  ▷ Estimate latency w/ sparsity pattern
6    $T_n^{Slack} \leftarrow SLO_n - Lat_n$  ▷ Calculate slack time
7    $Score_n \leftarrow Lat_n + \beta \times T_n^{Slack}$  ▷ Calculate score
8    $SendHW(Score_n, Info_n)$  ▷ Forward score & info to dynamic scheduler
9 end

```

representative requests offline. At the second level, the dynamic component interacts with the target hardware accelerator (neural processing unit (NPU) in Figure 8), to monitor the input-dependent, runtime sparsity of the workload. The obtained information is then used to update the task scores and adjust their processing order dynamically. We delve deeper into the details of *Dysta*'s static and dynamic components in the next subsection.

4.2 Dysta Scheduling Algorithm

4.2.1 Static Scheduler. Algorithm 1 describes the proposed software-based static scheduler. Upon the arrival of a new request $Reqst_n$, defined as a tuple $\langle Model_n, Pattn_n, input_n, SLO_n \rangle$, the static scheduler assigns an initial score $Score_n$ that determines its processing order before knowing any runtime information. To meet the multi-objective requirements of multi-DNN systems and maintain a balance between ANTT and SLO violation rate, we express the initial score of the n -th request as the weighted sum of two factors (line 7): the estimated latency Lat_n (line 5), and the slack time T_n^{Slack} that is determined based on the request's latency SLO (line 6).

Each of the two factors serves a different purpose. First, by incorporating the estimated latency in the score, we encourage shorter jobs to finish earlier. As such, it yields schedules that primarily improve the ANTT metric. Second, by integrating the slack into the score, our scheduler prioritizes jobs with tight deadlines to be dispatched sooner. Thus, the latency SLO violations are kept to a minimum. Overall, to balance the optimization between ANTT and SLO violation rate, we parametrize our formulation by means of hyperparameter β (line 7), which allows us to tunably weight the two factors.

Moreover, to capture the effect of sparsity pattern, latency is estimated separately for each model-pattern pair. We use the average latency to estimate Lat_n , which is obtained by collecting latency information either from actual executed requests or from using the performance model on a number of representative requests. A LUT is used to store the average latency and sparsity pattern pair of each model. As such, the estimated latency can be quickly obtained by accessing the LUT's entry for the given model-pattern pair. Finally, the calculated score and the request information are passed to the hardware-based dynamic scheduler for execution.

4.2.2 Dynamic Scheduler. Algorithm 2 presents the hardware-based dynamic scheduler. The dynamic scheduler uses a queue $ReqstQ$ to store all the incoming requests forwarded by the static scheduler. Our scheduling assumes that the execution is performed in a per-layer or per-layer-block manner, which is a common setting in commercial AI hardware [30, 48] and existing research literature [4, 5, 15, 43, 52]. Thus, whenever the execution of one layer or

Algorithm 2: Dynamic Scheduler

```

1 Input:  $Reqst_n = \langle Score_n, Info_n \rangle$  ▷ Input request from static scheduler
2 Output:  $Reqst_{next}$ 
3 if  $Reqst_n$  received then ▷ Request received, parallel execute with line 6
4    $Push(ReqstQ, Score_n, Info_n)$  ▷ Push request into queue
5 end
6 if  $LayerRun(Reqst_{cur})$  return at time  $t$  then ▷ Invoke scheduler
7    $\widehat{T}_{cur}^{Remain} \leftarrow PredSparseLat(S_{cur}^{Monitor})$  ▷ Update remain time
8   foreach  $Reqst_i$  in  $ReqstQ$  do ▷ Update per-request score
9      $T_i^{Slack} \leftarrow SLO_i - t - \widehat{T}_i^{Remain}$ 
10     $T_i^{Penalty} \leftarrow (T_i^{Wait} / T_i^{Isol}) / |ReqstQ|$ 
11     $Score_i \leftarrow \widehat{T}_i^{Remain} + \eta \times (T_i^{Slack} + T_i^{Penalty})$ 
12  end
13   $Reqst_{next} \leftarrow \underset{i}{\operatorname{argmin}} Score_i, \forall i \in [1, |ReqstQ|]$  ▷ Select new request
14 end

```

layer block completes, the dynamic scheduler is invoked to update the estimated latency and determine the next running request. As shown on line 7, the updated latency estimate, $\widehat{T}_{cur}^{Remain}$, is obtained based on the monitored sparsity information $S_{cur}^{Monitor}$ using an efficient sparse latency predictor, $PredSparseLat(\cdot)$, detailed in Section 5.1. To determine the next request to process, the dynamic scheduler updates the score of each $Reqst_i$ from $ReqstQ$, and the next running request is selected as the one with the lowest score.

The score calculated by the dynamic scheduler comprises three terms (line 9): 1) remaining time \widehat{T}_i^{Remain} , 2) slack time T_i^{Slack} , and 3) penalty $T_i^{Penalty}$. The first two terms aim to improve ANTT and violation rate, respectively, similarly to the static scheduler. However, a key characteristic of the dynamic component is that it estimates latency \widehat{T}_i^{Remain} with higher accuracy by relying on the actual monitored sparsity information and the sparse latency predictor. This approach allows us to alleviate the limitations demonstrated in Section 2.3.3, leading to better-informed scheduling decisions. To avoid excessively frequent preemptions, we introduce the penalty term (line 10). We define this term as the ratio between the waiting time T_i^{Wait} and the isolated execution time T_i^{Isol} , normalized by the number of requests in queue $ReqstQ$. A lower penalty indicates a request waiting for a shorter time, which encourages the scheduler to keep the nearest-executed request executing. A hyperparameter η is used to balance these terms in the score, allowing for a tunable trade-off between ANTT and SLO violation rate.

5 HARDWARE DESIGN

In this section, we focus on the hardware design of *Dysta*'s dynamic scheduler. We start by presenting the design of our sparse latency estimator, which constitutes a core driver behind the scheduler's decisions. Then, we proceed with the description of the microarchitecture, key components and hardware optimizations of our scheduler's design.

5.1 Sparse Latency Predictor

Designing the latency predictor for the sparse multi-DNN scheduler requires consideration of two metrics: accuracy and hardware overhead. As illustrated in Section 2.3.3, accurate latency estimation is crucial for improving the scheduling of sparse multi-DNN workloads. However, we must also take into account the hardware

Algorithm 3: Sparse Latency Prediction

```

1 Input:  $S^{\text{Monitor}}, i, j$                                 ▶ Monitored information from NPU
2 Output:  $Lat_i^{\text{Sparse}}$ 
3 if  $S^{\text{Monitor}}$  captured for  $j$ -th Layer of  $Reqst_i$  then    ▶ Monitored layer sparsity
4    $S_{(i,j)}^{\text{Avg}} \leftarrow \text{SparsityLUT}(i, j)$                 ▶ Get average sparsity
5    $Lat_i^{\text{Avg}} \leftarrow \text{LatencyLUT}(i)$                     ▶ Get average latency
6    $\gamma_i \leftarrow \text{PredSparsityCoeff}(S^{\text{Monitor}}, S_{(i,j)}^{\text{Avg}})$  ▶ Update sparsity coefficient
7    $Lat_i^{\text{Sparse}} \leftarrow \alpha \times \gamma_i \times Lat_i^{\text{Avg}}$     ▶ Predict latency
8 end
    
```




Figure 9: Correlation of sparsity of different layers.

overhead of the latency predictor. To this end, although several advanced learning-based approaches, such as Gaussian Processes [16], Random Forests [68] and DNNs [13], can be used for this task, the overhead of these methods is prohibitively costly for our scheduler operating at the layer granularity.

To design an accurate and efficient latency predictor for the sparse multi-DNN scheduler, two key problems need to be addressed: 1) identifying what sparsity information needs to be captured during runtime, and 2) determining the best algorithm for latency prediction. To this end, we profile the layer sparsity of two popular AttnN models, *BERT* and *GPT-2* on SQUAD [54] and GLUE [60] datasets, respectively, and analyse their Pearson product-moment correlation in Figure 9. The results indicate that the sparsities of different layers are highly linearly correlated in both models. This observation motivates us to monitor the layer sparsity at runtime and adopt a linear model for sparse latency prediction.

Algorithm 3 outlines our approach. We deploy a hardware monitor to calculate the layer sparsity S^{Monitor} for the current running model (line 1). The sparse latency predictor retrieves the average layer sparsity $S_{(i,j)}^{\text{Avg}}$ and latency Lat_i^{Avg} from the sparsity and latency LUTs, respectively, based on the model-pattern pair information (lines 4 & 5). The average latency stored in the latency LUT is received from the static scheduler, while the sparsity LUT is constructed by either collecting information from executed requests or obtaining information from model providers. Whenever the hardware monitor returns the layer sparsity, the sparsity coefficient γ_i is calculated (line 6), representing the linear rate between monitored and average layer sparsities. We then estimate the latency as $\alpha \times \gamma_i \times Lat_i^{\text{Avg}}$, where parameter α reflects how effectively sparsity can deliver real latency reduction. The value of α depends on the underlying hardware and needs to be set per pattern. As the accelerators targeted in this paper support both activation and weight sparsity, we set $\alpha = 1$.

Table 4: Root-Mean-Square Error (RMSE) of sparse latency predictor using three different strategies.

	<i>Average-All</i>	<i>Last-N</i>	<i>Last-One</i>
BERT	0.000286	0.000419	0.000252
GPT-2	0.000218	0.000421	0.000226




Figure 10: Overview of the proposed architecture.

To determine the sparsity coefficient γ , we consider three different approaches, namely *average-all*, *last-N*, and *last-one*. The *average-all* method takes the average of the monitored layer sparsity across all the already executed layers to estimate the dynamic layer sparsity. This estimated dynamic layer sparsity is then divided by the average sparsity obtained from the sparsity LUT to generate the sparsity coefficient. Similarly, *last-N* and *last-one* methods follow the same procedure as *average-all*, but estimate the dynamic layer sparsity from the monitored layer sparsity of the last N layers and the last executed layer, respectively. We evaluated these three approaches in the sparse latency predictor and compared the estimated latency with the measured sparsity obtained from the hardware simulator. Table 4 shows the root-mean-square error (RMSE) of latency prediction based on these three approaches, where parameter N of the *last-N* method is tuned through grid search and set to 3. We observe that both *average-all* and *last-one* perform similarly, and outperform *last-N*. As such, we opt to use the *last-one* method as it requires less amount of computation and memory for the averaging operation.

5.2 Hardware Scheduler Design

5.2.1 Design Overview. Figure 10 shows an overview of our hardware scheduler’s microarchitecture. The module is situated between the NPU and the host CPU, and is connected to the off-chip memory, if the intermediate results are transferred to it. We emphasize that the insertion point of our scheduler is dependent on the design of the hardware accelerator. As we do not impose any restrictions on the accelerator design, we utilize Figure 10 as an example to demonstrate the use of our hardware scheduler, without loss of generality.

The hardware scheduler consists of a controller, a runtime monitor, a compute unit and multiple FIFOs and LUTs. The FIFOs are deployed to track the per-request information, such as request tags, scores and latency SLOs. The depth of FIFOs depends on the maximal number of requests that can be handled by the hardware accelerator. We set the FIFO depths as configurable parameters in

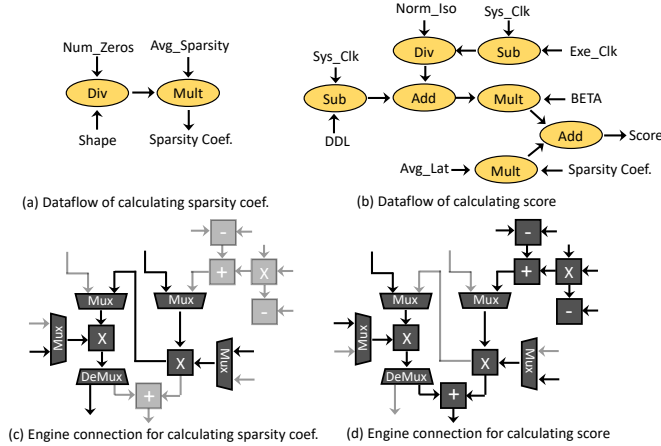


Figure 11: Compute engine with adaptable connections.

our scheduler. The latency, sparsity and shape information for each model-pattern pair are cached in three LUTs, which are accessed during the calculation of the sparsity coefficient and score. The controller is designed to perform the following tasks: 1) receive requests from the static scheduler and forward them into the tag and score FIFOs, 2) calculate the sparsity coefficient of the currently running request based on the monitored runtime information (Section 5.1), 3) update the score of each request, and 4) determine the request with the lowest score to be dispatched next. The calculation of the sparsity coefficient and the score are both implemented in the shared reconfigurable compute unit, detailed in Section 5.2.2. Finally, the monitor captures the layer sparsity through a zero-counting circuit.

5.2.2 Reconfigurable Compute Unit. To reduce the resource and area overheads of the hardware scheduler, we propose a reconfigurable compute unit that can be shared for the calculation of both the sparsity coefficient and the request score. As outlined in Algorithms 2 and 3, the computation comprises, first, calculating the sparsity coefficient γ and, then, updating the score of each request based on the estimated sparse latency. These computations are depicted as computational flows in Figure 11(a) and (b), respectively. As these computations are performed at different stages of the scheduling process, we design a reconfigurable compute unit that can be shared by both.

The hardware design of our reconfigurable compute unit is presented on the right-hand side of Figure 10. We equip the last two multipliers with several multiplexers (Mux) and a de-multiplexer (DeMux) at the input and output ports. This enables different dataflows to be implemented by reconfiguring the select signals. During the computation of the sparsity coefficient, as the shape is a constant, the division (Div) can be implemented as a multiplication by pre-computing the reciprocal of the shape offline. Therefore, the compute unit is configured with the last two multipliers enabled to calculate the sparsity coefficient, as shown in Figure 11(c). In contrast, during the computation of scores, all the arithmetic units are enabled for latency prediction and score aggregation, as shown in Figure 11(d). The division of normalized isolation time shown

Table 5: Comparison of scheduling approaches.

	Multi-AttNNs		Multi-CNNs	
	ANTT ↓	Violation Rate [%] ↓	ANTT ↓	Violation Rate [%] ↓
FCFS	18.9	55.1	11.4	23.1
SJF	5.0	15.2	2.6	3.4
SDRM ³ [33]	18.9	63.3	9.3	33.7
PREMA [7]	5.4	15.3	3.0	3.2
Planaria [17]	16.0	6.8	4.2	2.1
Dysta (Ours)	4.7	5.1	2.5	2.0

in Figure 11(b) is implemented as a multiplication by pre-computing the reciprocal offline. To further reduce resource consumption, we adopt half-precision floating-point (FP16) as the data type of the hardware scheduler.

6 EVALUATION

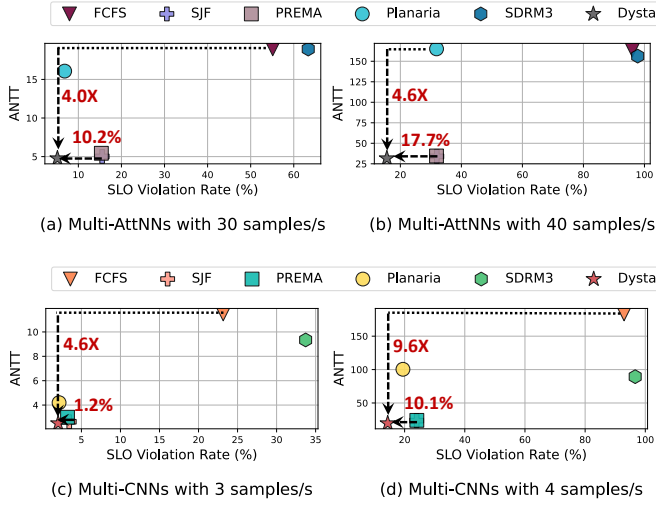
6.1 Experimental Setup

Software Implementation. We develop our evaluation and simulation infrastructure using *Python*, as described in Section 3.3. The pre-trained benchmark models (Table 3) are obtained from *TorchVision* provided by *PyTorch* (v3.8.12) and the *Transformers* library released by *HuggingFace*. The sparsification process described in Section 3.2 is implemented using *SparseML* [35]. We adopt the *Sanger* [43] simulator for sparse AttNNs experiments and a custom simulator of *Eyeriss-V2* for sparse CNNs. The custom *Eyeriss-V2* simulator is developed based on the validated performance model provided by a third-party code,[§] to ensure its correctness. The multi-AttNN workloads target the real scenarios on mobile phones with a mix of machine translation and question & answering as shown in Table 3. The Multi-CNNs constitute a mix of visual perception and hand tracking, which are common workloads in AR/VR [37], robotics [27] and data centers [59].

Hardware Implementation. We implement our hardware scheduler using *SystemVerilog*. The scheduler is clocked at 200 MHz. To make a consistent evaluation as *Eyeriss-V2* design[§] on FPGA, Xilinx Vivado 2019.1 is used for synthesis and implementation. To support large CNN models such as *ResNet-50* and *VGG-16*, we increase the global buffer of input activations in *Eyeriss-V2* from 1.5KB to 2.5KB. The other hardware configurations of both *Sanger* and *Eyeriss-V2* remain consistent with their original papers.

Baselines. We compare against status-quo scheduling approaches, including *i*) First-Come First-Served (FCFS) and *ii*) Shortest-Job First (SJF), and the state-of-the-art multi-DNN schedulers: *iii*) PREMA [7], *iv*) Planaria [17], and *v*) SDRM³ [33]. To improve the performance of PREMA at the beginning of the scheduling, we modify the selection criterion (line 9 of the PREMA scheduling algorithm presented in [7]) to $Token_i \geq Threshold$, as opposed to the original condition of $Token_i > Threshold$. For Planaria, we set the resource requirement estimate (line 40 in Algorithm 1 in [17]) to 1 for all tasks, since our target accelerators are time-shared, without spatial co-location

[§]https://github.com/SingularityKChen/dl_accelerator

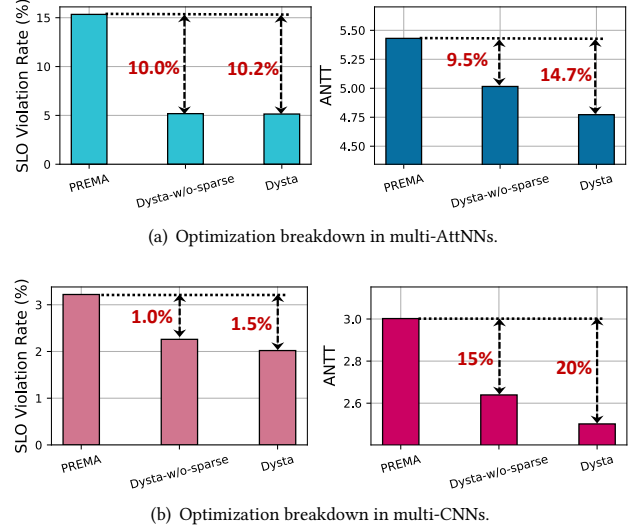

Figure 12: SLO violation rate and ANTT trade-off.

of multiple tasks. For SDRM³, we express MapScore (Eq. (5) in [33]) as the weighted sum of *Urgency* and *Fairness*. We further set the weight *Pref* to 1, since we target only one accelerator at a time, and tune parameter α following SDRM³'s optimization methodology.

Metrics. To assess the performance of different scheduling approaches, we utilize three metrics: average normalized turnaround time (ANTT), SLO violation rate and system throughput (STP). Specifically, for a multi-DNN workload consisting of N requests, we define ANTT as $\frac{1}{N} \sum_{n=1}^N \frac{T_i^{\text{Multi}}}{T_i^{\text{Isol}}}$ where T_i^{Multi} is the measured execution time under multi-tenancy and T_i^{Isol} is the uninterrupted isolated execution time of the target task. The SLO violation rate is calculated as $\frac{N^{\text{viol}}}{N}$ where N^{viol} represents the number of violated tasks under certain latency SLOs and request arrival rates. Following the experimental setup of [7], we set the SLO as $T_i^{\text{Isol}} \times M_{\text{slo}}^{\text{lat}}$ where $M_{\text{slo}}^{\text{lat}}$ refers to the SLO multiplier used to control the stringency of the latency constraints. We set the total number of requests in each workload as 1000 to ensure the reliability of our results. For each metric, we evaluate using five random seeds and report the average.

6.2 End-to-End Performance Comparison

We evaluate the performance of different scheduling approaches in two multi-tenant scenarios: multi-AttNNs and multi-CNNs. To generate the multi-AttNN and multi-CNN workloads, we randomly sample from AttNNs and CNNs listed in Table 3, respectively. Following the *MLPerf* standard [55], we generate the request arrival times for each workload using a Poisson distribution. Given the computational capacity of *Sanger* and *Eyeriss-V2*, we set the arrival rate of the multi-AttNN workload as 30 samples/s and the multi-CNN workload as 3 samples/s. The latency SLO multiplier is configured as 10 \times for both multi-AttNN and multi-CNN workloads. Table 5 presents the violation rate and ANTT of different scheduling approaches. We note that our proposed approach achieves


Figure 13: Optimization breakdown.

similar performance gains under workloads with different arrival rates and latency SLO multipliers, as demonstrated in Section 6.4.

As seen from Table 5, our *Dysta* scheduler outperforms all the previous approaches. In contrast to previous approaches, such as *PREMA* and *Planaria*, which only perform well on either violation rate or ANTT, our proposed scheduler excels in both metrics compared to the traditional heuristic *SJF* method. Comparing with the SOTA *Planaria*, our approach achieves a reduction of 1.7% in the violation rate while decreasing ANTT by 3.4 \times in multi-AttNN workloads. *Dysta* also reduces the ANTT by nearly 2 \times in multi-CNNs compared with *SjF* while achieving similar violation rates.

To visualize the improved trade-off achieved by *Dysta*, we present a 2D plot with ANTT on the y-axis and violation rate on the x-axis, as shown in Figure 12. The multi-AttNN workloads are evaluated at arrival rates of 30 and 40 samples/s, while the multi-CNN workloads are evaluated at arrival rates of 3 and 4 samples/s. The SLO configuration remains the same as in Table 5. As observed from Figure 12, *Dysta* is located at the lower left corner with the lowest violation rate and ANTT compared with other approaches, demonstrating its effectiveness in achieving an improved trade-off and better matching the multi-objective nature of multi-DNN systems.

6.3 Optimization Breakdown

To investigate the gain introduced by each proposed technique, we provide optimization breakdowns for both multi-AttNN and multi-CNN workloads, as shown in Figure 13. Two baselines are evaluated to compare with *Dysta*: 1) *PREMA*, a representative SOTA multi-DNN scheduling approach, and 2) *Dysta-w/o-sparse*, an ablated variant of *Dysta* with dynamic hardware scheduler and sparsity-aware support disabled. By comparing *Dysta-w/o-sparse* against *PREMA*, a clear reduction in the violation rate and ANTT can be observed on both multi-AttNN and multi-CNN workloads, demonstrating the effectiveness of our static score-based scheduling approach. Comparing *Dysta* to *Dysta-w/o-sparse*, we can see a clear ANTT drop on both multi-AttNN and multi-CNN workloads by adopting

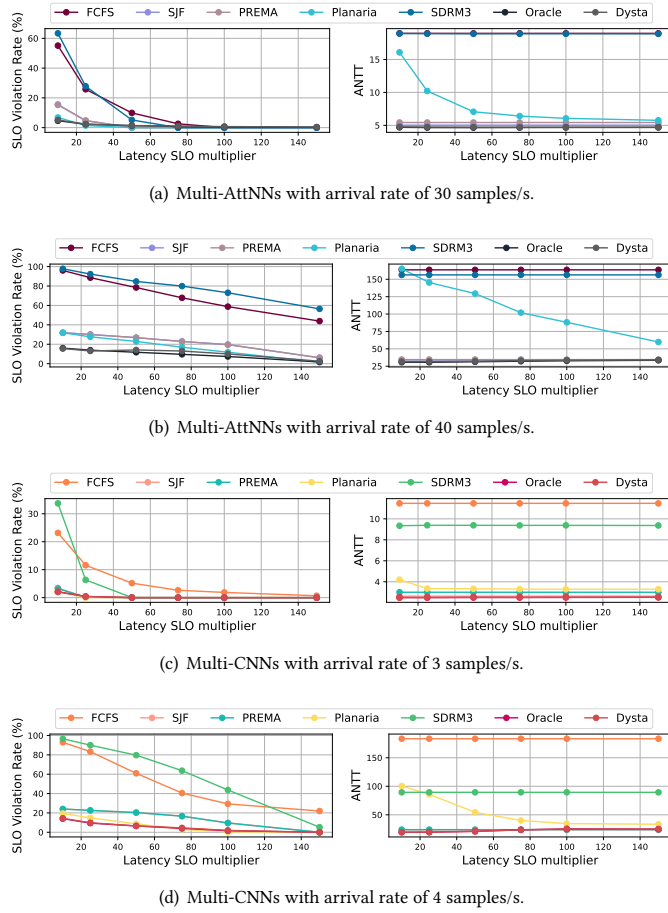


Figure 14: Evaluation across different latency SLOs.

the dynamic hardware component. In general, we observe that the dynamic hardware scheduler has a lesser impact on the violation rate than ANTT. This can be attributed to the significant correlation between task violations and latency SLO objectives. With loose SLO objectives, the latency estimated by the static scheduler is accurate enough to prevent task violation. Consequently, no further significant improvement is achieved by the dynamic hardware component.

6.4 Robustness Evaluation

To validate the robustness of our scheduling approach, we conduct experiments using workloads generated with different hyperparameters, *i.e.* latency SLOs and arrival rates.

6.4.1 Robustness across Different Latency SLOs. To evaluate the performance with different SLO requirements, we vary the latency SLO multiplier between $10\times$ and $150\times$ for both multi-AttNN and multi-CNN workloads. We evaluate two arrival rates for each workload, which are configured as 30 and 40 samples/s for multi-AttNNs and 3 and 4 samples/s for multi-CNNs. Figure 14 shows the results of the violation rate and ANTT. As the latency SLO multiplier increases, both ANTT and SLO violations exhibit a declining trend

due to the relaxed latency constraint. In multi-AttNN workloads, *Dysta* achieves the lowest violation rate and ANTT across different latency SLOs. This trend can be observed in both arrival rates (30 and 40 samples/s), demonstrating the advantage and robustness of our approach against different levels of latency SLO requirements. *Dysta* also closely matches the performance of the Oracle scheduler in ANTT and violation rate. The same conclusion can be drawn for multi-CNNs, where *Dysta* consistently outperforms other approaches in both violation rate and ANTT across different latency SLO multipliers. This also demonstrates the robustness of our approach under different families of models.

6.4.2 Robustness across Different Arrival Rates. To evaluate the performance across traffic levels, we vary the request arrival rate between 10 ~ 40 samples/s for multi-AttNNs, and 2 ~ 6 samples/s for multi-CNNs, based on the computational capacity of the target hardware. To eliminate the impact of latency SLOs, we set M_{slo}^{lat} to $10\times$ for both workloads. Figure 15 shows the violation rate, system throughput and ANTT of different scheduling methods across different arrival rates. All three metrics exhibit an upward trend as the arrival rate increases. The change in throughput remains the same across different scheduling methods as it is dependent on the computational capacity of the hardware. Under different arrival rates, *Dysta* keeps outperforming other schedulers in both violation rate and ANTT, with close performance to Oracle. The gain becomes more prominent while increasing the arrival rate, demonstrating the robustness of our method under heavier traffic.


6.5 Hardware Overhead

To evaluate the hardware overhead introduced by the proposed hardware scheduler, we compare its resource utilization against *Eyeriss-V2*. We adopt an open-source third-party hardware design[¶] to get the resource consumption of *Eyeriss-V2*. Both the hardware scheduler and accelerator are clocked at 200 MHz, targeting the Xilinx Zynq ZU7EV FPGA board.


To demonstrate the resource reduction provided by the reconfigurable compute unit and FP16 representation, we evaluate the resource usage of three hardware designs with different optimizations applied: 1) *Non_Opt_FP32*, which refers to the native implementation with separate compute units and 32-bit floating-point (FP32), 2) *Opt_FP32* that adopts reconfigurable compute unit, and 3) *Opt_FP16* with both reconfigurable compute unit and FP16 applied. To validate the effectiveness of our optimizations across different FIFO depths (Section 5.2.1), we instantiate two designs with FIFO depth of 512 and 64, respectively. As can be seen in Figure 16, the reconfigurable compute unit brings a significant reduction in LUT, register (Flip-Flops (FF)) and DSP resources. These reductions come from the savings of the compute unit, logic and memory required for calculating the sparsity coefficient. By comparing *Opt_FP16* with *Opt_FP32*, we can also see significant reductions in all three types of resources. For both optimizations, a similar reduction trend can be observed for both FIFO depths, demonstrating the effectiveness of our optimizations across different FIFO depths.

Table 6 presents the hardware overhead of our approach compared to *Eyeriss-V2*. Given the computational capacity of *Eyeriss-V2*,

[¶]https://github.com/karthisugumar/CSE240D-Hierarchical_Mesh_NoC-Eyeriss_v2



(a) SLO violation rate, throughput and ANTT of multi-AttNNs.



(b) SLO violation rate, throughput and ANTT of multi-CNNs.

Figure 15: Evaluation across different arrival rates on multi-AttNN and multi-CNN workloads.




Figure 16: Resource usage with different optimizations.

Table 6: Resource overhead of *Dysta* scheduler.

Module	LUTs	DSPs	On-Chip RAM
Eyeriss-V2	99168	194	140 KB
Scheduler	553	3	0.5 KB
<i>Dysta</i> -Eyeriss-V2	99721	196	140.5 KB
Total Overhead	0.55%	1.5%	0.35%

we set the FIFO depth to 64. Our hardware scheduler consumes a minimal amount of 0.55% of LUTs, 1.5% of DSPs and 0.35% of on-chip RAM, which is negligible compared to the overall resource usage. The significant gains achieved in ANTT and violation rate with such a low hardware cost demonstrates the effectiveness of our approach.

7 RELATED WORK

Multi-DNN Hardware Accelerators. Hardware architectures for multi-DNN workloads can be taxonomized based on their flexibility into: 1) fixed-purpose [36, 58], where the multi-DNN workloads are known *a priori*; and 2) workload-agnostic accelerators [2, 7, 17, 38]. *Fixed-purpose designs* opt to instantiate multiple heterogeneous compute engines [36, 58, 65] that operate in parallel and

apply design-time optimizations for the given multi-DNN workload. Prominent optimizations include highly customized compute engines and static scheduling [36, 58], both tailored to the target workloads, or co-design of the set of models and the hardware [65]. *Workload-agnostic designs* have either augmented existing accelerators with the necessary hardware components to support preemptive time-multiplexing [2, 7] or proposed highly flexible architectures that enable the spatial co-location of multiple DNNs through dynamic resource allocation [17, 38]. Finally, another line of work has investigated different trade-offs of multi-DNN hardware architectures through design space exploration [31, 32, 67] and interference-aware performance modeling [18]. In this work, we adopt preemptive time-multiplexing architectures, which constitute the most widespread deployed design. This allows *Dysta*'s techniques to improve performance with minimal hardware modifications, leading to wider generality and impact.

Multi-DNN Schedulers. Schedulers for multi-DNN workloads have so far adopted two main approaches: 1) temporal and 2) spatio-temporal scheduling. In the first approach, a single model occupies the compute engine of the target accelerator at each time instant. The optimizations have focused on preemptive scheduling schemes [7], inter-DNN pipelining [2, 49] that co-schedules compute- and memory-bound DNNs for higher utilization of both the computational and memory resources, and memory-aware scheduling to reduce memory swapping between models [8]. In the second approach, multiple models are processed in parallel by different parts of the accelerator. In this case, the scheduler decides both which DNNs to dispatch and the resource allocation among them, statically [31, 36, 58] or dynamically [17, 33, 38].

Despite the progress, as shown in Table 1, existing methods are constrained by two main limitations. First, by neglecting the presence of model sparsity, they leave untapped optimization opportunities. Second, each scheduler tends to work well on a single performance metric, e.g. either SLO violation rate or ANTT, while excessively penalizing the other. Instead, through its sparsity-aware,

fine-grained scheduling algorithm and the low-cost hardware enhancement, *Dysta* pushes the boundaries of the attainable performance, while consistently yielding a better trade-off between SLO violations and ANTT.

Multi-DNN Benchmarks. So far, progress has been made primarily towards benchmark suites for single-DNN execution, with prominent efforts such as AI Benchmark for mobile devices [25] and variants of *MLPerf* [26, 55]. Closer to our setting lies XRBench [37], a domain-specific benchmark suite of multi-DNN workloads for AR/VR applications. Nonetheless, there is still a lack of standardization in the evaluation of more generic multi-DNN systems, and their performance under sparse DNNs. We aim to pave the way towards bridging this gap through the open-source benchmark and evaluation infrastructure that we introduced in this work.

8 CONCLUSION

The increasing demand for running multiple DNNs in parallel and the prevalence of sparsity across different DNN models have led to the emergence of sparse multi-DNN workloads. By identifying the optimization opportunities in sparse multi-DNN workloads, we propose a novel bi-level dynamic and static scheduler that utilizes sparsity dynamicity and pattern information for better scheduling. Coupled with an efficient hardware scheduler and sparse latency predictor, our proposed approach achieves up to 10% fewer violations and nearly 4× lower average normalized turnaround time compared to the state-of-the-art methods, while incurring negligible hardware cost. To facilitate future development in this area, we will open-source all the benchmarks and code upon the paper acceptance. We believe our contributions will attract further research attention to study sparse multi-DNN workloads.

A ARTIFACT APPENDIX

A.1 Abstract

This Appendix summarizes the necessary information and instructions to reproduce our experimental results. Our artifacts mainly contain 1) public sparse multi-DNN benchmark, 2) simulation-based evaluation infrastructure for multi-DNN scheduling and 3) hardware prototype of the proposed dynamic scheduler. The profiling results and algorithmic performance can be reproduced by running our Python and PyTorch programs. The hardware resource consumption can be obtained by running Synthesis and Implementation on Vivado using our RTL code and constraint files. To facilitate reviewers and readers with quickly reproducing our results, we provide Bash scripts to re-generate all the figures and tables of our paper.

A.2 Artifact check-list (meta-information)

- **Algorithm:** *Dysta* scheduler, a sparsity-aware dynamic and static scheduling algorithm for sparse multi-DNN workloads.
- **Program:** Python, PyTorch, SparseML, Verilog HDL.
- **Model:** For multi-CNN workloads, we adopt four popular CNN models, namely *SSD* [41], *ResNet-50* [22], *VGG-16* [57], and *MobileNet* [24]. For multi-AttNNs workloads, our benchmark includes three commonly used language models: *BERT* [12], *BART* [39], and *GPT-2* [53].

- **Data set:** ImageNet [11], ExDark [42], DarkFace [63], and COCO [40] are used for vision tasks. GLUE [60] and SQUAD [54] are used for language tasks.
- **Run-time environment:** Ubuntu 20.04, CUDA SDK 11.3 or higher.
- **Hardware:** Nvidia RTX 2080 GPU, Intel Xeon Gold 6154 CPU.
- **Metrics:** Latency service-level objective (SLO) violation rate, average normalized turnaround time (ANTT) and hardware resource consumption.
- **Experiments:** Bash scripts and detailed instructions are provided to run experiments.
- **How much disk space required (approximately)?:** 20 GB.
- **How much time is needed to prepare workflow (approximately)?:** 1 ~ 2 hours.
- **How much time is needed to complete experiments (approximately)?:** Scheduling results need 3 hours. Resource consumption requires 5 hours.
- **Publicly available?:** Yes.
- **Code licenses (if publicly available)?:** Yes.
- **Archived (provide DOI)?:** 10.5281/zenodo.8252767

A.3 Description

A.3.1 How to access. You can access our codebase from the link: <https://github.com/SamsungLabs/Sparse-Multi-DNN-Scheduling>. Instructions on how to download the Docker image are provided in README.md.

A.3.2 Hardware dependencies. A GPU server is required to run the training of sparse models. A CPU server is needed to run simulation, synthesis and place-&-route.

A.3.3 Software dependencies. Vivado Design Suite 2019.2, PyTorch 1.11.0, CUDA SDK 11.3 or higher, Python 3.8 or higher. Other dependencies are described in README.md.

A.4 Installation

We provide a detailed installation guide in the README.md of the root directory.

A.5 Experiment workflow

To evaluate our artifacts, perform the following steps:

- Follow the instructions of experimental setup in the root directory to install software dependencies. Or Download the Docker image provided in Section A.3.1. Instructions on loading and creating the Docker container are provided in README.md.
- Download the necessary vision and language datasets. (Optional if using Docker)
- Generate CSV files using the hardware simulator in `hw_simulator`. (Optional if using Docker)
- Follow `INST_RPRODUCE.md` to run experiments using the scripts we provide.

A.6 Evaluation and expected results

We provide Bash scripts to generate all the figures and tables related to scheduling performance and resource consumption. Detailed

instructions on how to run scripts are described in `INST_RPROD UCE.md`. As running all the experiments requires a few hundred GPU/CPU hours, to facilitate the artifact evaluation, we refer to the following key results that can be obtained within a reasonable time:

- Profiling results of runtime variance while simulating different sparse models on CNN and AttNN accelerators. Following `INST_RPRODUCE.md` can reproduce Figure 2, Figure 3, Figure 4, and Table 2.
- ANTT and latency SLO violation rate trade-off as shown in Table 5, Figure 12 and Figure 13. Following `INST_RPROD UCE.md` to run Bash scripts under folder `dysta_scheduler/script`.
- Stress test of scheduling approaches under different latency SLOs and arrival rates as shown in Figure 14 and Figure 15. Following `INST_RPRODUCE.md` to run Bash scripts under folder `dysta_scheduler/script`.
- Hardware resource consumption under different hardware optimizations. Following `INST_RPRODUCE.md` to download Vivado reports and run the script under `/hw_design/draw_fig_hw_opt.sh`

A.7 Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-badging>
- <http://cTuning.org/ae/submission-20201122.html>
- <http://cTuning.org/ae/reviewing-20201122.html>

REFERENCES

- [1] Mario Almeida, Stefanos Laskaridis, Abhinav Mehrotra, Lukasz Dudziak, Ilias Leontiadis, and Nicholas D. Lane. 2021. Smart at What Cost? Characterising Mobile Deep Neural Networks in the Wild. In *ACM Internet Measurement Conference (IMC)*.
- [2] Eunjin Baek, Dongup Kwon, and Jangwoo Kim. 2020. A Multi-Neural Network Acceleration Architecture. In *International Symposium on Computer Architecture (ISCA)*.
- [3] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. 2020. What is the State of Neural Network Pruning?. In *Conference on Machine Learning and Systems (MLSys)*.
- [4] Yu-Hsin Chen et al. 2016. Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks. *IEEE Journal of Solid-State Circuits (JSSC)* 52, 1 (2016), 127–138.
- [5] Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, and Vivienne Sze. 2019. Eyeriss v2: A Flexible Accelerator for Emerging Deep Neural Networks on Mobile Devices. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS)* 9, 2 (2019), 292–308.
- [6] Seungbeom Choi, Sunho Lee, Yeonjae Kim, Jongse Park, Youngjin Kwon, and Jaehyuk Huh. 2022. Serving Heterogeneous Machine Learning Models on Multi-GPU Servers with Spatio-Temporal Sharing. In *USENIX Annual Technical Conference (ATC)*.
- [7] Yujeong Choi and Minsoo Rhu. 2020. PREMA: A Predictive Multi-task Scheduling Algorithm for Preemptible Neural Processing Units. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [8] Bart Cox, Jeroen Galjaard, Amirasoud Ghiassi, Robert Birke, and Lydia Y Chen. 2021. Masa: Responsive Multi-DNN Inference on the Edge. In *IEEE International Conference on Pervasive Computing (PerCom)*.
- [9] Shailesh Dave, Riyadh Baghdadi, Tony Nowatzki, Sasikanth Avancha, Aviral Shrivastava, and Baoxin Li. 2021. Hardware Acceleration of Sparse and Irregular Tensor Computations of ML Models: A Survey and Insights. *Proc. IEEE* 109, 10 (2021), 1706–1752.
- [10] Chunhua Deng, Yang Sui, Siyu Liao, Xuehai Qian, and Bo Yuan. 2021. GoSPA: An Energy-efficient High-performance Globally Optimized SParse Convolutional Neural Network Accelerator. In *Annual International Symposium on Computer Architecture (ISCA)*.
- [11] Jia Deng et al. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *ACL*.
- [13] Lukasz Dudziak, Thomas Chau, Mohamed Abdelfattah, Royson Lee, Hyeji Kim, and Nicholas Lane. 2020. BRP-NAS: Prediction-based NAS using GCNs. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [14] Stijn Eyerman and Lieven Eeckhout. 2008. System-level Performance Metrics for Multiprogram Workloads. *IEEE MICRO* 28, 3 (2008), 42–53.
- [15] Hongxiang Fan, Thomas Chau, Stylianos I Venieris, Royson Lee, Alexandros Kouris, Wayne Luk, Nicholas D Lane, and Mohamed S Abdelfattah. 2022. Adaptable Butterfly Accelerator for Attention-based NNs via Hardware and Algorithm Co-design. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 599–615.
- [16] Martin Ferianc, Hongxiang Fan, Divyansh Manocha, Hongyu Zhou, Shuanglong Liu, Xinyu Niu, and Wayne Luk. 2021. Improving Performance Estimation for Design Space Exploration for Convolutional Neural Network Accelerators. *Electronics* 10, 4 (2021), 520.
- [17] Soroush Ghodrati, Byung Hoon Ahn, Joon Kyung Kim, Sean Kinzer, Brahendra Reddy Yatham, Navateja Alla, Hardik Sharma, Mohammad Alian, Eiman Ebrahimi, Nam Sung Kim, et al. 2020. Planaria: Dynamic Architecture Fission for Spatial Multi-Tenant Acceleration of Deep Neural Networks. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*.
- [18] Shikha Goel, Rajesh Kedia, M Balakrishnan, and Rijurekha Sen. 2020. INFER: INterference-aware Estimation of Runtime for Concurrent CNN Execution on DPUs. In *IEEE International Conference on Field-Programmable Technology (ICFPT)*.
- [19] Tae Jun Ham, Sung Jun Jung, Seonghak Kim, Young H Oh, Yeonhong Park, Yoonho Song, Jung-Hun Park, Sanghee Lee, Kyoung Park, Jae W Lee, et al. 2020. A³: Accelerating Attention Mechanisms in Neural Networks with Approximation. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [20] Tae Jun Ham, Yejin Lee, Seong Hoon Seo, Soosung Kim, Hyunji Choi, Sung Jun Jung, and Jae W Lee. 2021. ELSA: Hardware-Software Co-Design for Efficient, Lightweight Self-Attention Mechanism in Neural Networks. In *ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*.
- [21] Song Han, Huizi Mao, and William J Dally. 2015. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In *International Conference on Representation Learning (ICLR)*.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [23] Yihui He, Xiangyu Zhang, and Jian Sun. 2017. Channel Pruning for Accelerating Very Deep Neural Networks. In *IEEE International Conference on Computer Vision (ICCV)*. 1389–1397.
- [24] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv* (2017).
- [25] Andrey Ignatov, Radu Timofte, Andrei Kulik, Seungsoo Yang, Ke Wang, Felix Baum, Max Wu, Lirong Xu, and Luc Van Gool. [n. d.]. AI Benchmark: All About Deep Learning on Smartphones in 2019. In *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*.
- [26] Vijay Janapa Reddi, David Kanter, Peter Mattson, Jared Duke, Thai Nguyen, Ramesh Chukka, Ken Shiring, Koan-Sin Tan, Mark Charlebois, William Chou, et al. 2022. MLPerf Mobile Inference Benchmark: An Industry-Standard Open-Source Machine Learning Benchmark for On-Device AI. In *Conference on Machine Learning and Systems (MLSys)*.
- [27] Hanhwi Jang, Joonsung Kim, Jae-Eon Jo, Jaewon Lee, and Jangwoo Kim. 2019. Mnfnst: A fast and scalable system architecture for memory-augmented neural networks. In *46th International Symposium on Computer Architecture (ISCA)*. 250–263.
- [28] Jun-Woo Jang, Sehwan Lee, Dongyoung Kim, Hyunsun Park, Ali Shafiee Ardestani, Yeongjae Choi, Channoh Kim, Yoojin Kim, Hyeongseok Yu, Hamzah Abdel-Aziz, et al. 2021. Sparsity-Aware and Re-configurable NPU Architecture for Samsung Flagship Mobile SoC. In *ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*.
- [29] Joo Seong Jeong, Jingyu Lee, Donghyun Kim, Changmin Jeon, Changjin Jeong, Youngki Lee, and Byung-Gon Chun. 2022. Band: Coordinated Multi-DNN Inference on Heterogeneous Mobile Processors. In *Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*.
- [30] Norman P Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, et al. 2023. TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings. In *ACM/IEEE 50th Annual International Symposium on Computer Architecture (ISCA)*.

- [31] Sheng-Chun Kao and Tushar Krishna. 2022. MAGMA: An Optimization Framework for Mapping Multiple DNNs on Multiple Accelerator Cores. In *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*.
- [32] Rajesh Kedia, Shikha Goel, M. Balakrishnan, Kolin Paul, and Rijurekha Sen. 2020. Design Space Exploration of FPGA-Based System with Multiple DNN Accelerators. *IEEE Embedded Systems Letters (ESL)* (2020).
- [33] Seah Kim, Hyoukjun Kwon, Jinook Song, Jihyuck Jo, Yu-Hsin Chen, Liangzhen Lai, and Vikas Chandra. 2024. SDRM³: A Dynamic Scheduler for Dynamic Real-time Multi-model ML Workloads. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [34] Mark Kurtz, Justin Kopinsky, Rati Gelashvili, Alexander Matveev, John Carr, Michael Goin, William Leiserson, Sage Moore, Nir Shavit, and Dan Alistarh. 2020. Inducing and Exploiting Activation Sparsity for Fast Inference on Deep Neural Networks. In *International Conference on Machine Learning (ICML)*.
- [35] Mark Kurtz, Justin Kopinsky, Rati Gelashvili, Alexander Matveev, John Carr, Michael Goin, William Leiserson, Sage Moore, Nir Shavit, and Dan Alistarh. 2020. Inducing and Exploiting Activation Sparsity for Fast Inference on Deep Neural Networks. In *International Conference on Machine Learning (ICML)*. 5533–5543.
- [36] Hyoukjun Kwon, Liangzhen Lai, Michael Pellauer, Tushar Krishna, Yu-Hsin Chen, and Vikas Chandra. 2021. Heterogeneous Dataflow Accelerators for Multi-DNN Workloads. In *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*.
- [37] Hyoukjun Kwon, Krishnakumar Nair, Jamin Seo, Jason Yik, Debabrata Mohapatra, Dongyuan Zhan, Jinook Song, Peter Capak, Peizhao Zhang, Peter Vajda, et al. 2023. XRBench: An Extended Realistic (XR) Machine Learning Benchmark Suite for the Metaverse. In *Conference on Machine Learning and Systems (MLSys)*.
- [38] Jounghoo Lee, Jinwoo Choi, Jaeyoon Kim, Jinho Lee, and Youngsook Kim. 2021. Dataflow Mirroring: Architectural Support for Highly Efficient Fine-Grained Spatial Multitasking on Systolic-Array NPU. In *Design Automation Conference (DAC)*.
- [39] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*.
- [40] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision (ECCV)*. Springer, 740–755.
- [41] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. SSD: Single Shot Multibox Detector. In *European Conference on Computer Vision (ECCV)*. Springer, 21–37.
- [42] Yuen Peng Loh and Chee Seng Chan. 2019. Getting to Know Low-Light Images with the Exclusively Dark Dataset. *Computer Vision and Image Understanding* 178 (2019), 30–42.
- [43] Liqiang Lu, Yicheng Jin, Hangrui Bi, Zizhang Luo, Peng Li, Tao Wang, and Yun Liang. 2021. Sanger: A Co-Design Framework for Enabling Sparse Attention Using Reconfigurable Architecture. In *54th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*.
- [44] Gaurav Menghani. [n. d.]. Efficient Deep Learning: A Survey on Making Deep Learning Models Smaller, Faster, and Better. *ACM Computing Surveys (CSUR)* ([n. d.]).
- [45] Francisco Muñoz Martínez, Raveesh Garg, Michael Pellauer, José L. Abellán, Manuel E. Acacio, and Tushar Krishna. 2023. Flexagon: A Multi-Dataflow Sparse-Sparse Matrix Multiplication Accelerator for Efficient DNN Processing. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [46] Francisco Muñoz-Martínez, José L. Abellán, Manuel E. Acacio, and Tushar Krishna. 2021. STONNE: Enabling Cycle-Level Microarchitectural Simulation for DNN Inference Accelerators. In *IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 201–213.
- [47] Vinod Nair and Geoffrey E Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *International Conference on Machine Learning (ICML)*.
- [48] NVIDIA. 2021. Accelerating Inference with Sparsity using Ampere and TensorRT. <https://developer.nvidia.com/blog/accelerating-inference-with-sparsity-using-ampere-and-tensorrt/>. Accessed: April 7, 2024.
- [49] Young H Oh, Seonghak Kim, Yunho Jin, Sam Son, Jonghyun Bae, Jongsung Lee, Yeonhong Park, Dong Uk Kim, Tae Jun Ham, and Jae W Lee. 2021. Layerweaver: Maximizing Resource Utilization of Neural Processing Units via Layer-Wise Scheduling. In *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*.
- [50] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [51] Eric Qin, Ananda Samajdar, Hyoukjun Kwon, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and Tushar Krishna. 2020. SIGMA: A Sparse and Irregular GEMM Accelerator with Flexible Interconnects for DNN Training. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [52] Zheng Qu, Liu Liu, Fengbin Tu, Zhaodong Chen, Yufei Ding, and Yuan Xie. 2022. DOTA: Detect and Omit Weak Attention for Scalable Transformer Acceleration. In *27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [53] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [54] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2383–2392.
- [55] Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chou, Ramesh Chukka, Cody Coleman, Sam Davis, Pan Deng, Greg Diamos, Jared Duke, Dave Fick, J. Scott Gardner, Itay Hubara, Sachin Idgunji, Thomas B. Jablin, Jeff Jiao, Tom St. John, Pankaj Kanwar, David Lee, Jeffery Liao, Anton Lohmotov, Francisco Massa, Peng Meng, Paulius Mickevicius, Colin Osborne, Gennady Pekhimenko, Arun Tejusve Raghunath Rajan, Dilip Sequeira, Ashish Sirasao, Fei Sun, Hanlin Tang, Michael Thomson, Frank Wei, Ephrem Wu, Lingjie Xu, Koichi Yamada, Bing Yu, George Yuan, Aaron Zhong, Peizhao Zhang, and Yuchen Zhou. 2020. MLPerf Inference Benchmark. In *International Symposium on Computer Architecture (ISCA)*.
- [56] Francisco Romero, Qian Li, Neeraja J. Yadwadkar, and Christos Kozyrakis. 2021. INFaaS: Automated Model-less Inference Serving. In *USENIX Annual Technical Conference (ATC)*.
- [57] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations (ICLR)*.
- [58] S. I. Venieris and C. S. Bouganis. 2018. f-CNN^x: A Toolflow for Mapping Multiple Convolutional Neural Networks on FPGAs. In *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*.
- [59] Stylianos I Venieris, Christos-Savvas Bouganis, and Nicholas D Lane. 2023. Multi-DNN Accelerators for Next-Generation AI Systems. *Computer* 56, 3 (2023), 70–79.
- [60] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (2018).
- [61] Erwei Wang, James J Davis, Ruizhe Zhao, Ho-Cheung Ng, Xinyu Niu, Wayne Luk, Peter YK Cheung, and George A Constantinides. 2019. Deep Neural Network Approximation for Custom Hardware: Where We've Been, Where We're Going. *ACM Computing Surveys (CSUR)* 52, 2 (2019), 1–39.
- [62] Hanrui Wang, Zhekai Zhang, and Song Han. 2021. SpAtten: Efficient Sparse Attention Architecture with Cascade Token and Head Pruning. *IEEE International Symposium on High Performance Computer Architecture (HPCA)* (2021).
- [63] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. 2018. Deep Retinex Decomposition for Low-Light Enhancement. In *British Machine Vision Conference (BMVC)*.
- [64] Yannan Nellie Wu, Po-An Tsai, Angshuman Parashar, Vivienne Sze, and Joel S Emer. 2022. Sparseloop: An Analytical Approach To Sparse Tensor Accelerator Modeling. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*.
- [65] Lei Yang et al. 2020. Co-Exploration of Neural Architectures and Heterogeneous ASIC Accelerator Designs Targeting Multiple Tasks. In *Design Automation Conference (DAC)*.
- [66] Juheon Yi and Youngki Lee. 2020. Heimdall: Mobile GPU Coordination Platform for Augmented Reality Applications. In *Annual International Conference on Mobile Computing and Networking (MobiCom)*.
- [67] Shulin Zeng, Guohao Dai, Niansong Zhang, Xinhao Yang, Haoyu Zhang, Zhenhua Zhu, Huazhong Yang, and Yu Wang. 2022. Serving Multi-DNN Workloads on FPGAs: a Coordinated Architecture, Scheduling, and Mapping Perspective. *IEEE Transactions on Computers (TC)* (2022).
- [68] Li Lyna Zhang, Shihao Han, Jianyu Wei, Ningxin Zheng, Ting Cao, Yuqing Yang, and Yunxin Liu. 2021. nn-Meter: Towards Accurate Latency Prediction of Deep-Learning Model Inference on Diverse Edge Devices. In *International Conference on Mobile Systems, Applications, and Services (MobiSys)*.
- [69] Aojun Zhou, Yukun Ma, Junnan Zhu, Jianbo Liu, Zhijie Zhang, Kun Yuan, Wenxiu Sun, and Hongsheng Li. 2021. Learning N: M Fine-grained Structured Sparse Neural Networks from Scratch. In *International Conference on Representation Learning (ICLR)*.
- [70] Zhe Zhou, Junlin Liu, Zhenyu Gu, and Guangyu Sun. 2021. Energon: Towards Efficient Acceleration of Transformers Using Dynamic Sparse Attention. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* (2021).