

Graft: Efficient Inference Serving for Hybrid Deep Learning with SLO Guarantees via DNN Re-alignment

Jing Wu, Lin Wang, *Senior Member, IEEE*, Qirui Jin, Fangming Liu*, *Senior Member, IEEE*

Abstract—Deep neural networks (DNNs) have been widely adopted for various mobile inference tasks, yet their ever-increasing computational demands are hindering their deployment on resource-constrained mobile devices. Hybrid deep learning partitions a DNN into two parts and deploys them across the mobile device and a server, aiming to reduce inference latency or prolong battery life of mobile devices. However, such partitioning produces (non-uniform) DNN fragments which are hard to serve efficiently on the server. This paper presents Graft—an efficient inference serving system for hybrid deep learning with latency service-level objective (SLO) guarantees. Our main insight is to mitigate the non-uniformity by a core concept called DNN re-alignment, allowing multiple heterogeneous DNN fragments to be restructured to share layers. To fully exploit the potential of DNN re-alignment, Graft employs fine-grained GPU resource sharing. Based on that, we propose efficient algorithms for merging, grouping, and re-aligning DNN fragments to maximize request batching opportunities, minimizing resource consumption while guaranteeing the inference latency SLO. We implement a Graft prototype and perform extensive experiments with five types of widely used DNNs and real-world network traces. Our results show that Graft improves resource efficiency by up to 70% compared with the state-of-the-art inference serving systems.

Index Terms—Deep Learning Systems, Edge Computing, Hybrid Deep Learning, GPU Sharing.

1 INTRODUCTION

With the rapid development of deep learning (DL), we have seen a variety of DL-based mobile applications such as augmented reality [1], [2], [3] and intelligent personal assistant [4]. These applications typically employ pre-trained deep neural networks (DNNs) to perform inference tasks such as object detection and natural language understanding. While achieving unprecedentedly high accuracy, modern DNNs are bloated in size and impose excessive computational demands [5], [6]. Meanwhile, mobile inference requires low latency, typically specified with a service-level objective (SLO), to ensure good user experience [7], [8], [9], [10], [11]. Recent advances in mobile accelerators and DNN optimizations mitigate this mismatch issue, but bring concerns over battery life or incur extra costs from model retraining [12], [13], [14], [15].




Fig. 1: Hybrid deep learning across the mobile device and the edge server, resulting in misaligned DNN fragments on the edge server, which are hard to serve efficiently.

Hybrid DL tackles this challenge by extending the mobile device with servers typically in an edge cloud [9], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25]. Specifically, hybrid DL partitions a DNN into two parts (referred to as fragments) and executes them on the mobile device and the server, respectively, as depicted in Figure 1. Unlike server-only solutions where the DNN is offloaded completely to a server, hybrid DL is still able to exploit the computing resources available on the mobile device for partial DNN execution and can accommodate network dynamics through adaptive DNN partitioning [23]. Hybrid DL has been widely explored in both academia [9], [23], [26], [27] and industry [16], [18] for minimizing the inference latency or maximizing the battery life of mobile device. It has been shown that hybrid DL can reduce the inference latency of popular DNNs by up to $1.69\times$ [28]. In addition, hybrid DLs

- This work was supported in part by National Key Research & Development (R&D) Plan under grant 2022YFB4501703, and the Major Key Project of PCL under Grant PCL2022A05. Lin Wang was supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 210487104 - SFB 1053. (Corresponding author: Fangming Liu)
- J. Wu and Q. Jin are with the National Engineering Research Center for Big Data Technology and System, the Services Computing Technology and System Lab, Cluster and Grid Computing Lab in the School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan 430074, China. E-mail: wujinghust@hust.edu.cn and qiruijin@umich.edu.
- L. Wang is with Paderborn University and TU Darmstadt, Germany. E-mail: lin.wang@uni-paderborn.de.
- F. Liu is with Peng Cheng Laboratory, and Huazhong University of Science and Technology, China. E-mail: fangminghk@gmail.com.

offer enhanced privacy protection since they only exchange intermediate data, which blurs a significant level of personal information due to local inferences [11], [29], [30].

Despite the vast successes, efficiently serving the generated DNN fragments on the server side remains a critical challenge, which has been overlooked so far. In particular, different mobile devices running the same DNN may produce varying partition decisions due to their different yet changing network conditions. Consequently, the server will need to handle a large number and variety of *misaligned* DNN fragments with relatively small inference request rates from individual mobile devices, as depicted in Figure 1. Although these DNN fragments share layers partially, the lack of uniformity due to the misalignment largely limits the opportunity of applying request batching—an essential technique for improving resource efficiency in virtually all inference serving systems [10]. Existing inference serving systems are either agnostic to such partial sharing and would simply treat these DNN fragments separately, or blindly merge architecture-identical fragments without further considering other individual requirements (e.g., time budget) for servers, leading to poor resource efficiency that is detrimental to edge clouds.

This paper addresses this timely challenge by introducing the concept of *DNN re-alignment* for hybrid DL. The key idea is to re-partition, i.e., apply to another partitioning, non-uniform DNN fragments such that we can generate aligned DNN fragments that consist of the same DNN layers. These aligned DNN fragments can then be served by shared DNN instances where request batching can be applied as in existing DL inference serving systems. By deciding the re-partition point carefully, we can leverage request batching to its full potential, thus improving the server resource efficiency in hybrid DL. The entire optimizations on non-uniform fragments for higher resource efficiency refer to re-alignment.

We propose Graft—a first-of-its-kind inference serving system adopting DNN re-alignment for hybrid DL. *The goal of Graft is to achieve resource efficiency while guaranteeing latency SLO.* To this end, Graft re-aligns the DNN fragments on the server to enable request batching across inference requests received from different mobile devices. Graft features efficient strategies for merging and grouping DNN fragments by converting the grouping problem into a variant of the balanced graph partitioning problem. Graft also incorporates a fine-grained resource allocation policy based on spatial GPU sharing enabled by CUDA MPS (Multi-Process Service) of NVIDIA GPUs. Specifically, Graft formulates joint re-alignment and resource allocation as an optimization problem and produces near-optimal decisions on the re-partition point, batch size, GPU allocation per instance, and the number of instances for all fragments on the server. Ultimately, Graft minimizes the GPU resource consumption on the server while ensuring that the end-to-end latency for each inference request is bounded by the SLO.

Overall, the paper makes the following contributions: after identifying the challenges of provisioning misaligned DNN fragments in hybrid DL (§2), we

- introduce a new concept called DNN re-alignment

and present the design of Graft—a first-of-its-kind inference serving system tailored for hybrid DL with latency SLO guarantees, employing DNN re-alignment (§3).

- present efficient algorithms for DNN re-alignment, featuring greedy DNN fragment merging, graph-partitioning based fragment grouping, and fragment re-partitioning and resource allocation, with near-optimal resource efficiency (§4).
- implement Graft and perform extensive experiments with five different DNNs to evaluate Graft based on a system prototype (§5). Experiment results show that Graft is able to achieve resource savings by up to 70% when compared with the state-of-the-art inference serving systems, while guaranteeing latency SLOs.

§7 discusses related work and §8 draws the conclusions.

2 BACKGROUND AND MOTIVATION

We present the background and discuss the challenges and insights to motivate our work in this section.

2.1 Deep Learning Deployment

Deep learning (DL) has become the de facto approach for inference tasks in domains like computer vision and natural language processing [4], [23], [31]. To improve accuracy, sophisticated DNNs with a large number of parameters have been explored [32], requiring substantial computing resources to execute. For example, the computing demand (GFLOPS) of image classification DNNs has increased by up to $595.6\times$ since 2016 [33] and the memory requirement of DNNs for high-resolution image transformations has reached 6-8GB [14]. This poses critical challenges for DL deployment, especially on resource-constrained mobile devices [23], [34], [35].

One solution is to fit the DNN execution on mobile devices via DNN optimizations including pruning and quantization [36], [37], [38], [39], [40], [41], [42], [43]. Without optimizations, a DNN for image segmentation can drain the 4000mAh battery of a mobile device by 95.9% in just one hour when running on mobile CPUs [12]. DNN optimizations reduce the computational demands but suffer from an accuracy gap and lack practical deployment support. For example, empirical studies show that only 10.3% of the DNN-based apps from the Google Play store employ DNN optimizations [12]. Developers tend to build customized DNNs for their apps with lower resource consumption but at the cost of lower accuracy [44]. The associated costs for optimizing and customizing a DNN for diverse mobile devices are considerable [36], [45]. Meanwhile, over 50% of the existing mobile devices still use processors that are at least six years old [46], which clearly falls short of meeting the resource requirements of modern DNNs in entirety.

Another solution is to offload the DNN execution completely to a powerful server equipped with accelerators like GPUs, typically at the network edge over a wireless network [47], [48], [49]. This approach avoids the needs for DNN optimizations thanks to the high capability of the server, but it brings two other issues (as illustrated in Figure 2 (top)): (1) Leaving the resources on mobile devices




Fig. 2: Changing resource consumption (top) and DNN partition points (middle) of Inception-v3 in hybrid DL under a snippet (50s) of a real-world 5G network trace [55].

unused leads to excessive resource consumption on the server. (2) The inference latency suffers from uncertainty due to the network variability, leading to potentially high SLO violations.

2.2 Hybrid Deep Learning

To overcome these limitations, *hybrid DL* aims to leverage computing resources on both the mobile device and the edge server for DL inference, without sacrificing on accuracy. Considering the layered structure of DNNs, the key idea is to partition the DNN into two fragments at an intermediate layer, with each fragment mapped to the mobile device and to the server, respectively [9], [23], [50], [51], [52]. Hybrid DL has been widely explored by the academic (e.g., Neurosurgeon [23] and SPINN [9]) and the industry (e.g., Taobao [16], [17] and Huawei Cloud [18], [53]), with different partitioning strategies proposed to minimize the inference latency or maximize the battery life of the mobile device. Recently, it has been shown that hybrid DL can achieve a speedup of $1.69\times$ and reduce mobile energy consumption by 22.5% on average [28], [54].

Since the network bandwidth is likely to be dynamic, DNN partitioning thus needs to be adapted continuously. Figure 2 (middle) shows the partition point variation under a real-world 5G network bandwidth trace (Figure 2(bottom)) [55] when applying hybrid DL. We can see from Figure 2(top) that hybrid DL accommodates the variability of network bandwidth; meanwhile, by exploiting the resources on mobile devices, hybrid DL can reduce the server resource consumption by up to $3\times$ and avoids SLO violations via dynamic DNN partitioning. So far, the focus of existing work has been on the partitioning strategy and little attention has been paid to the efficient serving of the DNN fragments on the edge server.

2.3 DL Inference Serving for Hybrid DL

The proliferation of DL-based applications renders the critical importance of DL inference serving. There exist a large




Fig. 3: Throughput (RPS) achievable vs. throughput demanded in hybrid DL with GPUs allocated by Clipper [7]. Here, sharing enables a single instance to be used when the two mobile clients experience the same network condition and thus the same partition point.




Fig. 4: DNN re-alignment via re-partitioning.

variety of DL inference serving systems, such as TensorFlow Serving [56], Clipper [7], Nexus [8], Clockwork [10], ALERT [57], and INFaaS [6], etc., that concern scheduling DL inference requests to appropriate resources (e.g., GPUs), meeting SLOs on inference latency with the least amount of resources.

Despite all these efforts, virtually all inference serving systems suffer efficiency issues when applied to *hybrid DL* due to the following reasons: (1) Dynamic DNN partitioning in hybrid DL produces *non-uniform* DNN fragments that cannot be served with shared DNN instances on servers. Therefore, a separate DNN instance may need to be provisioned for each individual mobile client with a small request rate, leading to limited request batching opportunities critical for resource efficiency. Also, these DNN fragments change continuously due to network dynamics of the mobile clients. (2) Existing systems mostly employ coarse-grained GPU allocation [7], [10], where DNNs are encapsulated in containers that occupy GPUs exclusively in a time-sharing manner [58]. This leads to considerable GPU under-utilization [8], [58]. Figure 3 depicts the throughput achievable by the allocated resources versus the actual demand from two mobile clients each issuing 30 requests per second (RPS). With Clipper the allocated GPU resources can support more than an order of magnitude higher request rates, even with opportunistic sharing (a single instance is used when the two mobile clients have the same network condition and thus the same partition point) enabled between the two mobile clients.

2.4 DNN Re-alignment and Challenges

To address the non-uniformity issue, we observe that a simple technique called *DNN re-alignment* could help. The key idea is to re-partition DNN fragments to generate aligned DNN fragments shared by multiple mobile clients, as depicted in Figure 4. This works since DNN fragments




Fig. 5: Discreteness in resource consumption of Inception-v3: (a) required GPU share to meet different inference time budgets with throughput of 200RPS, (b) required GPU share to achieve different inference throughput while maintaining inference latency of 25ms.

in hybrid DL, despite being misaligned, share their last layers. This sharing behavior can even be observed across different DNNs due to the widespread adoption of pre-trained foundation models and transfer learning (for fine-tuning the final layers) [12], [58]. Re-alignment leverages such sharing to increase the request batching opportunity, aiming to achieve better resource efficiency.

However, DNN re-alignment amplifies DNN fragmentation, which, when combined with coarse-grained GPU allocation, leads to poor GPU utilization. Fortunately, the state-of-the-art CUDA MPS (Multi-Process Service) allows for sharing a GPU spatially across multiple processes to improve utilization. Based on MPS, recent work like GSLICE with fine-grained GPU allocation demonstrates significant resource savings for inference serving [59]. Therefore, we adopt MPS and use GPU share (%) to represent the amount of GPU resources.

Efficiently applying DNN re-alignment for inference serving in hybrid DL is non-trivial. In particular, we identify the following three major challenges.

How to merge DNN fragments? There may exist DNN fragments that are uniform in structure, but serve requests with different time budgets and/or different request rates. Merging the requests for these fragments and using one instance to serve them could potentially increase the batching opportunity, but the time budget of all requests will need to follow the smallest one among all merged requests. This hurts resource efficiency since some more relaxed time budget will be wasted. Interestingly, we observe that smaller time budgets and higher throughput do not always translate into higher resource consumption, as shown in Figure 5. This is caused by the discreteness of the batch size and the resource unit. This property can be leveraged to merge DNN fragments to improve resource efficiency.

How to group DNN fragments for re-alignment? Given a set of non-uniform DNN fragments, deciding the subset of the DNN fragments to be grouped for re-alignment is critical. DNN fragments can be heterogeneous with respect to multiple factors: the partition point, the time budget, and the request rate. There are two decisions to make for fragment grouping, namely, how many DNN fragments to put in one group and which DNN fragments to put together? Intuitively, a larger group size provides more opportunities for optimization, but comes with exponentially higher time




Fig. 6: System architecture of Graft.

complexity for making grouping decisions. The grouping decision should be made based on the compound influence of the heterogeneity factors on resource efficiency, which needs to be explored carefully.

How to select the re-partition point? After grouping, we need to perform re-alignment by selecting a layer to re-partition the DNN fragments in the same group. This re-partition point can be at any layer between the partition point of the largest fragment and the last layer. Hence, deciding the re-partition point consuming minimum resources requires to explore all the possible layers efficiently.

3 GRAFT SYSTEM DESIGN

We present Graft—a new system for serving DNNs fragments in hybrid DL. The goal of Graft is to minimize the overall server resource consumption while guaranteeing latency SLOs for mobile applications.

Overview. The system architecture of Graft is depicted in Figure 6. Graft consists of three core components: *profiler*, *scheduler*, and *executor*, all running on the edge server. The profiler runs offline while the scheduler and the executor run online at runtime. We assume each mobile client applies hybrid DL by partitioning the DNN with any existing strategies and we focus on serving the DNN fragments with GPUs on the server side. The general workflow of Graft is as follows: First, the profiler collects performance profiles of DNNs with respect to both latency and throughput in batch size and GPU share across all available GPUs in servers. This ensures Graft is adaptable to heterogeneous GPU setups. Then, the scheduler takes the performance profiles and generates a resource-efficient execution plan. The execution plan consists of the DNN fragment groups, the re-partition point for each fragment in each group, the GPU share for each fragment, and the number of instances to spawn for each fragment. Finally, the executor takes the execution plan and deploys the DNN fragment instances with specified GPU shares accordingly. Here, the requests for each DNN fragment are evenly distributed to all its instances with a load balancer. Additionally, the requests that fail to meet SLOs are dropped by load balancer for resource saving.

Due to effective optimizations (detailed in § 4.3), Graft maintains high time-efficiency in its decision-making. This offers the opportunity to repeatedly invoke the scheduler when any mobile device changes its partition point due to significant network bandwidth changes. This trigger-based approach adjusts its resource allocation timely based

on real-time network conditions, significantly boosting the inference serving performance.

Scheduler. The scheduler provides the intelligence of the system by generating the execution plan, taking as input the offloaded DNN fragments, the request rate of each mobile client, and the performance profiles for the concerned DNNs from the profiler. The goal of scheduling is to minimize the overall resource consumption for the given mobile clients, while guaranteeing their latency SLOs. To this end, the scheduler first applies an efficient algorithm to merge fragments (§4.1), and then organizes DNN fragments into groups (§4.2). Finally, it performs a fine-grained resource allocation algorithm with pruning-based optimizations (§4.3) to determine the most resource-efficient execution plan, which is then deployed by the executor to serve the mobile clients.

4 SCHEDULING FOR DNN RE-ALIGNMENT

We now discuss the three steps the scheduler takes to generate the most resource-efficient execution plan, namely fragments merging, grouping, and re-partitioning.

4.1 DNN Fragments Merging

Resource allocation for serving DNNs involves decision-making for multiple discrete variables including the number of DNN instances, unit of GPU share, and batch size. Such discreteness can lead to significant resource over-allocation when provisioning a large number of DNN fragments, i.e., a DNN instance can serve more requests without requiring more resources. We quantify resource over-allocation with a metric called *resource margin* defined as $(q_a - q_d)/q_d$, which captures the gap between the demanded throughput q_d and the actual achievable throughput q_a under given resources.

Reducing resource over-allocation is equivalent to minimizing the resource margin. To this end, *we propose to merge incrementally uniform fragments that have the same partition point and time budget, until the resource margin decreases to a predefined threshold called the merging threshold.* Merging with such a threshold leaves more space for optimization in fragments grouping and re-partitioning, which turns out to be preferable for Graft over the strategy where we merge as much as possible. We will analyze the impact of the merging threshold in §5.5. Merging not only improves resource efficiency, but also reduces the number of fragments to deal with in the next steps of grouping and re-partitioning.

4.2 DNN Fragments Grouping

The second step concerns the grouping of DNN fragments on the server into fragment subsets such that the total resource consumption by the DNN fragments is minimized. Obtaining the optimal grouping of fragments is hard, due to the combinatorial explosion. One insight we can leverage is that *it is beneficial to group fragments with similar properties (i.e., partition point, time budget, and request rate).* Applying this insight, we show that *the grouping problem can be converted into a variant of the classic balanced graph partitioning problem* [60], which aims to divide the vertices of a graph into K equal-sized and disjoint subsets while minimizing/maximizing the total edge costs across different subsets.

Intuitively, we can construct a complete graph with all DNN fragments as nodes and for the edges between node pairs, we assign weights based on the similarity of the fragments represented by the nodes on each edge. Now, the grouping problem becomes dividing the nodes into K equal-sized and disjoint subsets on the constructed graph. The goal is to maximize the total edge weights in all subsets, equivalent to minimizing the total edge weights across subsets.

Solving the balanced graph partition problem is still non-trivial [60]. Existing work typically applies heuristics that either maximize internal edge costs within subsets or minimize external edge costs. Fennel et al. propose a framework that combines heuristics in both directions [61]. We follow a similar approach and define an objective function integrating the variance of the edge weights in each subset and the total edge weights across subsets, formulated as

$$\min \sum_{k=1}^K \sum_{e \in \mathbf{E}_k} \frac{(w_e - \bar{w}_k)^2}{|\mathbf{E}_k|} + \sum_{k=1}^K \sum_{e \in \mathbf{E}'_k} w_e, \quad (1)$$

where w_e represents the weight on edge e , \bar{w}_k is the average edge weight in subset k , \mathbf{E}_k and \mathbf{E}'_k are the sets of internal and external edges for subset k , respectively. We assign the edge weight using the weighted Euclidean distance between the property vectors (consisting of partition point p , time budget t , and inference throughput q) of the fragments on each edge, i.e., $|\vec{v}_i, \vec{v}_j|$ where $\vec{v} = \langle p, t, q \rangle$.

To strike the trade-off between resource efficiency and time complexity, we design a greedy algorithm following the approach in [61] with the above objective function: (a) Randomly select K fragments to initiate the groups. Here, K is the number of expected groups and is dictated by the group size. Note that the group size does not constrain the batch size, which is decided in resource allocation during re-partitioning. (b) Iterate over the remaining fragments and assign each of them to the group with the least increase of cost given by Formula (1). We will explore the impact of group size and factor weights in the Euclidean distance calculation in §5.6.

4.3 DNN Fragments Re-partitioning

We now focus on deciding the re-partition point and resource allocation for the fragments in each group, aiming to minimize the overall resource consumption with latency SLO guarantees. *Re-partitioning breaks down the fragment execution into two stages (see Figure 4): the alignment stage where fragments have their own parts to execute concurrently, and the shared execution stage where the same instances are used for all fragments.* Given a re-partition point, deciding the resource allocation is equivalent to allocating appropriate time budget to each of the two stages, since the time budget dictates the required resources in each stage.

The re-alignment algorithm is listed in Algorithm 1. Given a group \mathbf{F}_G of fragments, we loop through all possible re-partition points (line 5) between $\min\{p_1, \dots, p_M\}$ and p_E , where $p_i, i \in [1, M]$ denotes the partition point of fragment f_i in the group and p_E denotes the last layer of the DNN. Under each re-partition point p , we can further divide the fragments into two parts: \mathbf{F}_A that will be re-partitioned and \mathbf{F}_B that will not be affected (line 7). For \mathbf{F}_A we apply

Algorithm 1: Fragments re-partitioning algorithm

Input: $F_G = \{(p, t, q)\}$: a group of fragments
Output: $\langle p_{OPT}, X_{OPT} \rangle$: min-resource re-partition point and corresponding execution plan

```

1 Function realign( $F$ ):
2   if  $|\mathbf{F}| = 0$  then
3     return  $\emptyset$ 
4    $\mathbf{W} \leftarrow \emptyset$ 
5   foreach  $p \in [\min\{p_1, \dots, p_M\}, p_E]$  do
6      $r_{min} \leftarrow \infty, X \leftarrow \emptyset$ 
7      $\mathbf{F}_A \leftarrow \{f_i \mid f_i \in \mathbf{F} \wedge p_i < p\}, \mathbf{F}_B \leftarrow \mathbf{F} \setminus \mathbf{F}_A$ 
8      $\mathbf{D} \leftarrow \{\langle d_1, \dots, d_{|\mathbf{F}_A|+1} \rangle \mid d_i + d_{|\mathbf{F}_A|+1} \leq \min\{t_j \mid f_j \in \mathbf{F}_A\}/2 \wedge i \in [1, |\mathbf{F}_A|]\}$ 
9     foreach  $\mathbf{d} \in \mathbf{D}$  do
10       $\langle \mathbf{r}, \mathbf{b}, \mathbf{c} \rangle \leftarrow \text{min\_resource}(p, \mathbf{d})$ 
11      if  $\sum \mathbf{r} \leq r_{min}$  then
12         $r_{min} \leftarrow \sum \mathbf{r} \circ \mathbf{c}, X \leftarrow \langle \mathbf{r}, \mathbf{b}, \mathbf{c} \rangle$ 
13       $\mathbf{w} \leftarrow \{p, X\} \cup \text{realign}(\mathbf{F}_B)$ 
14      return  $\mathbf{W} \cup \{\mathbf{w}\}$ 
15  $\mathbf{S} \leftarrow \text{realign}(F_G)$ 
16  $\langle p_{OPT}, X_{OPT} \rangle \leftarrow \arg \min_{\mathbf{w} \in \mathbf{S}} \sum \mathbf{r} \circ \mathbf{c}$ 
17 return  $\langle p_{OPT}, X_{OPT} \rangle$ 

```

the re-partitioning at layer p and obtain all possible time budget allocation schemes denoted by \mathbf{D} , where each allocation scheme is given by $\mathbf{d} = \langle d_1, \dots, d_{|\mathbf{F}_A|+1} \rangle$ and d_i is the time allocated to fragment f_i . Here, $d_{|\mathbf{F}_A|+1}$ represents the time allocated for the common fragment shared among all fragments in \mathbf{F}_A generated by the re-partitioning. A valid time budget allocation scheme has to ensure that the total time spent by a request is lower than the available time budget of that request. *Considering the worst-case queuing delay which equals the execution time [8], the total execution time of each fragment after re-partitioning has to be no larger than half of the total available time budget (line 8).* Meanwhile, batch sizes (denoted as \mathbf{b}) are self-adapted with respect to re-partition points and time budgets, rather than fixed settings. For each legitimate scheme $\mathbf{d} \in \mathbf{D}$, we calculate the minimum required resources and find out the scheme with the smallest resource consumption (lines 9-12). The above procedure will be recursively applied to the fragment set \mathbf{F}_B until no fragment is left (line 13). The final step is to find the re-partition point and the corresponding min-resource execution plan that leads to the lowest overall resource consumption (line 16).

The time complexity of re-alignment algorithm can be analyzed as follows. The merging step uses “mergesort” to sort the fragments in $O(n \log n)$ time where n is the total number of fragments. The grouping step is transformed into a variant of balanced graph partitioning, which is solved by our proposed greedy method in $O(n^2)$ time. The re-partitioning step loops through $O(n/M)$ groups, and in each group the re-partitioning problem (concerning M fragments) can be solved by the Simplex algorithm in polynomial time [62]. We will explore the time complexity of Graft empirically in §5.9.

Graft applies several important optimizations to reduce

TABLE 1: Specifications of mobile devices.

Device	GPU	AI perf.	Mem.	Mode
Nano	128-core Maxwell	472 GFLOPS	4G	MAXN [63]
TX2	256-core Pascal	1.33 TFLOPS	8G	MAXQ [64]

its time overhead: (1) DNN fragments merging (§4.1) to reduce the number of fragments, (2) graph building based on a simple metric and heuristics to group fragments with an adjustable group size (§4.2), and (3) pruning search space by reserving only the most “efficient” solutions (e.g., the blue dots in Figure 5a) and parallelism for re-partitioning for different groups.

5 EVALUATION

We have implemented a prototype of Graft and conducted extensive experiments to validate its performance.

5.1 Setup and Implementation

We assess the effectiveness of Graft across three different experimental scales, namely, small-, large- and massive scales, which are categorized based on the number of involved mobile devices. In small-scale experiments, four NVIDIA Jetson Nano Developer Kits and two NVIDIA Jetson TX2 Developer Kits are used as mobile devices to evaluate Graft’s performance under homogeneous/heterogeneous mobile scenarios. The specifications of mobile devices are given in Table 1. Following existing work [9], [26], we choose Jetson boards since they have comparable specifications to real-world low- and high-end mobile devices, but are much simpler to program and control. In large-scale experiments, we follow the widely adopted approach where we use two additional servers (in total 20 CPU cores) to emulate 20 mobile devices [9], [59]. Additionally, we also use simulation to evaluate Graft under massive mobile devices (e.g., at the level of thousands). All servers and Jetson boards are connected to the same switch with 10GbE and 1GbE links, respectively.

The data path of the system is implemented as follows. Each DNN fragment is materialized as multiple DNN instances, each of which works as an independent process, by DNN runtime Pytorch v1.9.0. The DNN is partitioned across the mobile device and the edge server based on Neurosurgeon [23] for its simplicity; other partitioning strategies can be used as well (explained in §9). As Figure 6 shows, a request, once completes on-device processing, is first sent, via a network socket, to the server and then buffered as a tensor in a queue for the corresponding fragment. This queue is shared by all the instances for each DNN fragment, which process requests in batch from the queue. At server-sides, Graft utilizes the Unix domain socket to transfer intermediate result (tensors) among DNN fragments. Moreover, we set up a thread pool with a capacity of five for each DNN instance to efficiently pipeline their transferring.

The control path implements the scheduling algorithms in Python. For implementing the merging and grouping algorithms, we leverage packages `numpy` and `networkx`. For calculating the optimal re-partition point and resource allocation, i.e., execution plan, we use `cvxpy` with the `GUROBI`

TABLE 2: Parameters and inference latency of used models when executing on mobile GPUs with batch size as one and executing on server GPUs with GPU share as thirty and batch sizes as one.

Model	Inc	Res	VGG	Mob	ViT
Number of layers	17	16	6	18	15
Mobile latency [Nano] (ms)	165	226	147	84	816
Mobile latency [TX2] (ms)	94	114	77	67	603
Server latency (ms)	29	30	6	19	58

solver, which takes the profiles of DNN fragments with respect to both latency and throughput in batch size and GPU share. Based on the plan, Graft starts DNN instances accordingly while terminating the outdated instances in an asynchronous manner. The network bandwidth between the mobile device and the server is shaped by replaying real-world 5G bandwidth traces [55] (trace snippet shown in Figure 2), using the Linux `tc` utility (detailed in §9).

We use five representative DNNs from TorchVision v0.10.0: four image classification models namely Inception-v3 (Inc), ResNet-101 (Res), VGG11 (VGG), and Vision Transformer-B16 (ViT), and a semantic image segmentation model called DeepLabV3 MobileNetv3-Large (Mob) [31]. The input size for the DNNs is around 588KB [59]. DNNs, except for ViT, have request rates of 30RPS from each mobile device. Due to its high mobile latency (column 6 in Table 2), ViT has a lower request rate of 1RPS to prevent excessive frames blocking on mobile sides, which could deteriorate user experience. Table 2 lists the parameters and inference latency of these models. The latency SLO for the considered models is typically lower than their mobile inference latency; otherwise, there is no need to offload to a server. We set the SLO of each model to 95% of the model’s mobile inference latency. Later, we will explore the performance of Graft under varying SLOs in §5.10.

The focus of Graft is on inference serving on the edge server. Therefore, we use the state-of-the-art inference serving system GSLICE [59] and its enhanced version GSLICE⁺, the optimal allocation (denoted as “Optimal”), and the static allocation (Static) and its enhanced version (Static⁺) as our baselines. Note that Static and Static⁺ depend on the average bandwidth of each mobile device to decide its resource allocation. Neither GSLICE/GSLICE⁺ nor Static/Static⁺ conducts DNN re-alignment.

We evaluate two performance metrics namely the total resource consumption and the end-to-end latency (i.e., SLO guarantees). Meanwhile, for a fair comparison, the baselines follow the same hybrid DL setup as Graft. Specifically, requests’ statistics, namely initial partition points and time budgets, over the small- and the large-scale scenarios, are depicted in Figure 7a and Figure 7b, respectively. Note that Res, Mob, and ViT exhibit relatively polarised partitioning, with fragments displaying high similarities in initial partition points. This is attributed to their unique architectures, which enables significant reductions in transmission loads at specific layers. Taking Mob as an instance, layer one achieves reductions by 71.1% compared with raw input. Consequently, Neurosurgeon makes its polarised decisions to reap lower latency.

The experiments involving the Optimal, namely §5.2,

TABLE 3: Overall resource reduction by Graft compared with GSLICE and GSLICE⁺ under small- and large-scale, respectively.

	Model	Inc	Res	VGG	Mob	ViT
Small-scale	Homo. (%)	35.8	24.2	19.1	25	70.0
	Heter. (%)	24.2	16.3	31	26.4	63.1
Large-scale	Homo. (%)	30.5	5.8*	16.5	29.2	11.6*
	Heter. (%)	20.4	16.1	30.3	41.1	59.6

* Case where GPU memory is the bottleneck.

§5.3, §5.6, §5.9 and §5.10, are repeated by 10 times considering their considerable time cost; meanwhile, the remaining experiments are conducted by 50 times, unless specified otherwise. The resource allocation for each fragment on GPU is enforced via MPS with the resource unit of 1% GPU share. To reduce the interference among concurrent instances, we empirically cap the allocated GPU shares of each GPU to be lower than 100% [59].

5.2 Performance in Small-Scale Scenarios

Homogeneous cases. We compare the *resource consumption* of Graft, GSLICE, Optimal, and Static when serving the models. Here, we use four Jetson Nanos as mobile devices. Table 3 (row 1) shows the average resource reduction of Graft when compared with GSLICE and Figure 8a depicts the detailed comparison. We can see that *Graft outperforms GSLICE and Static significantly and performs close to Optimal. This improvement is attributed to DNN re-alignment, which allows for a larger batch size for each fragment, thus improving the resource efficiency.*

Figure 9 shows the *end-to-end latency* distribution. We observe that both Static and GSLICE outperform Graft with Inc, Res, VGG, and ViT, due to its resource over-allocation. Notably, Graft achieves lower latency with Mob. This is because Mob has relatively low heterogeneity in both partition points and time budgets in the small-scale scenario (depicted in Figure 7a and Figure 7b), allowing Graft to re-align the four fragments simultaneously. Moreover, considering SLO guaranteeing, Graft depends on the lowest time budget of the fragments to make its decision, which speeds up the execution of some fragments than they demand, thus achieving lower latency.

Heterogeneous cases. We now compare the *resource consumption* of Graft, GSLICE, Optimal, and Static under the heterogeneous scenarios, with four Jetson Nanos and two Jetson TX2s as mobile clients. Table 3 (row 2) shows the average resource reduction of Graft when compared with GSLICE; details are given in Figure 8b. Meanwhile, Graft achieves resource reduction by 13.9%, 21%, 33.9%, 28.8%, and 63.1% for Inc, Res, VGG, Mob, and ViT, respectively when compared with Static. We can observe from Figure 7b that Jetson TX2 generates fragments with significantly lower time budgets than those from Jetson Nanos; meanwhile, the difference in partitioning points of DNN fragments also magnifies as illustrated in Figure 7a. Facing this heterogeneous scenario, Graft tends to re-align the fragments separately based on their mobile devices, for maintaining high resource efficiency without violating SLOs. This prevents Graft from building large batches across heterogeneous

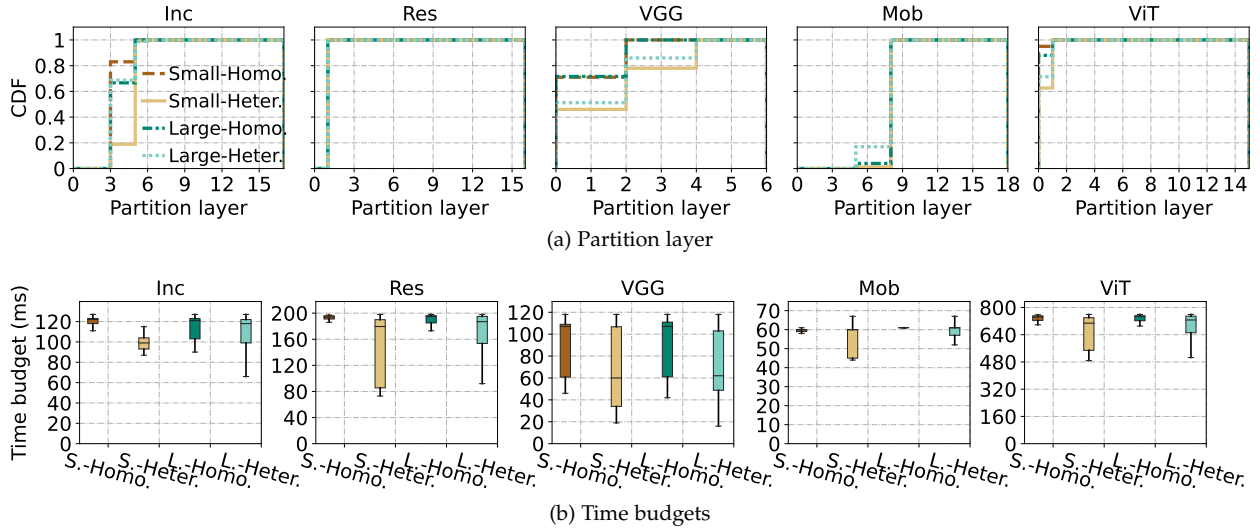


Fig. 7: Initial partition layer and time budgets of DNN fragment. S./L. represents small and large scales respectively.




Fig. 8: Resource consumption comparison in (a) small-scale scenarios with four Nano boards, (b) small-scale scenarios with four Nano and two TX2 boards, and (c) large-scale scenarios with 20 emulated mobile clients.

devices, especially for DNNs with limited resource margins (e.g., Res), thus affecting its effectiveness. On the contrary, large gains are observed with VGG. This is attributed to its high resource margins, which accommodate variations in time budgets and allow for larger batches that span across heterogeneous devices.

We depict the *end-to-end latency* distribution of models deployed on Nano and TX2 in Figure 10a and Figure 10b, respectively. We observe that Graft meets latency SLOs in most cases, except the case with VGG on Jetson TX2. The reason is that VGG re-aligns the fragments from the heterogeneous mobile devices, thus forming large batches. While improving the resource efficiency, the increased batch size leads to the variation in queueing delay and incurs long-tail latency, especially for TX2 which has tight SLOs.

5.3 Performance in Large-Scale Scenarios

Homogeneous cases. We compare the *resource consumption* of Graft, GSLICE, Optimal, and Static under 20 mobile devices (emulated by 20 CPU cores). The increased number of mobile devices introduces uniform fragments that share the same initial partition point and time budget. For a fair comparison, we enable the best merging strategy (i.e., merge all uniform fragments) for GSLICE and Static and denote the enhanced version as GSLICE⁺ and Static⁺. To prevent GPU memory overflow, we set an upper bound for the number of instances to five for each fragment based on

empirical observations. This bound is removed in large-scale simulations in §5.8.

Table 3 (row 3) presents the average resource reduction of Graft when compared with GSLICE⁺; details are in Figure 8c. We can see that *Graft consumes significantly fewer resources than GSLICE⁺, and the gap between Graft and Optimal is within 4%*. Notably, the constraints in the number of instances for each fragment narrow down the optimization space, degrading the performance of Graft. In particular, Res and ViT are bottlenecked by their significant requirements in GPU memory, beyond our testbed’s capacity, and get diminished improvements in resource efficiency. This constraint also leads to resource over-allocation for Res as 63% and 37.1% of Graft than GSLICE and Static⁺, respectively. Overall, *Graft reduces the resource consumption by 36.1%, 20.8%, 46.9%, and 74.7% for Inc, VGG, Mob, and ViT, respectively, when compared with GSLICE. The reduction is 32.9%, 9.1%, 25.7%, and 11.6% as compared with Static⁺.*

Figure 11 depicts the *end-to-end latency* distribution where Graft provides SLO guaranteeing for the five DNNs. Note that GSLICE⁺ provides lower end-to-end latency due to its considerable resource over-allocation when compared with Graft and Optimal.

Heterogeneous cases. We also evaluate Graft in large-scale heterogeneous scenarios, where we emulate 15 Nano and 5 TX2 boards as mobile devices. Applying the algorithms of Graft, we calculate the resource consumption and compare






Fig. 9: End-to-end latency distribution comparison in small-scale scenarios with four Nano boards. (The vertical dashed line marks the latency SLO).



(a) Jetson Nano



(b) Jetson TX2

Fig. 10: End-to-end latency distribution comparison in small-scale heterogeneous scenarios for (a) Nano and (b) TX2 boards.




Fig. 11: End-to-end latency distribution comparison in large-scale scenarios with 20 emulated mobile clients. (The vertical dashed line marks the latency SLO).

it with that of $GSLICE^+$. As shown in Table 3 (row 4), even when provisioning DNNs with polarised partitioning, namely Res, Mob, and ViT (explained in §5.1), which are particularly suitable for $GSLICE^+$ to merge and batch due to their similar initial partition points, Graft still achieves lower resource consumption by 16.1%, 41.1%, and 59.6%, respectively. This is attributed to Graft’s comprehensive exploration in re-partition points and fine-grained re-allocation of time budgets (detailed in §4.3), which constructs substantial optimization spaces for improving resource efficiency.

We were not able to obtain the end-to-end latency distribution due to the lack of GPU memory to support that number of DNN fragments under strict SLOs.

5.4 Effectiveness of Re-partitioning

We now evaluate the resource consumption with and without re-partitioning when provisioning five random fragments, irrespective of the impact of merging and grouping. Here, each fragment has the same frame rate as in §5.2 and replays a random bandwidth from the real-world network trace. Figure 12 shows that *re-alignment reduces the resource consumption by up to 32%, 10%, 15.6%, 26.7% and 60% for*

Inc, Res, VGG, Mob, and ViT, respectively. The improvement for Res is smaller due to its relatively lower resource margin.




Fig. 12: Resource consumption with re-partitioning, normalized by the case without re-partitioning.

We also analyze how well re-partitioning performs under different network bandwidths and request rates. To this end, we evaluate the re-partition point and the required GPU share while keeping four fragments unchanged and varying the configuration for the fifth fragment. Taking Inc as an example, we observe from Figure 13a that *the resource consumption decreases marginally with the increase of the network bandwidth*. The reason can be seen in Algorithm 1 that re-alignment is dictated by the lowest time budget of fragments in the same group. With the increase of network bandwidth, the partition point of the fragments is fixed at

layer three. However, their time budget difference keeps increasing until the time budget of the static fragments becomes the bottleneck for re-alignment, resulting in a marginal decrease in resource consumption. Additionally, we observe in Figure 13b that *high request rates lead to higher resource consumption*. As for the re-partition point, it is influenced by the DNN type, partition point, time budget, and request rate, simultaneously.




Fig. 13: Re-partition point and GPU share of Inc under varying (a) network bandwidth and (b) request rate.

5.5 Effectiveness of Merging

We evaluate the resource consumption of Graft when employing three different merging strategies: No-merging, Uniform merging (i.e., merging all uniform fragments), and Uniform⁺ merging (i.e., merging until reaching the merging threshold), respectively. We generate 50 fragments and set the merging threshold as 0.2. Figure 14 shows that, *compared with Uniform⁺ merging, Uniform merging suffers resource over-allocation when applied to Res*. This occurs because Res fragments exhibit relatively low heterogeneity as shown in Figure 7a, which leads to a number of fragments with high throughput when conducting Uniform merging. Moreover, Res has constrained resource margins due to its high computation intensity. The conflict, between high throughput and low resource margin, often leads to the independent execution of DNN fragments with no further realignment, which hurts the resource efficiency consequently. In comparison, *Uniform⁺ merging limits the number of fragments involved in merging and controls the fragments throughput. This creates reasonable space for realignment across fragments, thus achieving higher resource efficiency. Furthermore, this improvement grows with the increase of fragment numbers as shown in Figure 15(top).*




Fig. 14: Resource consumption with different merging strategies (50 fragments with merging threshold 0.2).

It is worth noting that *insufficient merging, e.g., adopting No merging or setting Uniform⁺ with high merging threshold,*

also hurts resource efficiency. This is because it results in limited available batch sizes for realignment. On the other hand, *Uniform⁺ merging greatly decreases the problem sizes for the following grouping and re-partitioning*. We can see from Figure 15(bottom) that Uniform⁺ merging reduces the number of DNN fragments by up to 42.6%, 42.3%, 50.7%, 76.1%, and 38.9% on average for Inc, Res, VGG, Mob, and ViT respectively, as compared with No-merging.




Fig. 15: Resource consumption (top) and time cost (bottom) of Res normalized by no merging under different numbers of DNN fragments with merging threshold of 0.2.

We analyze the sensitivity of Uniform⁺ merging to different merging thresholds under varying fragment numbers. As illustrated in Figure 16a, *most DNNs except Res are insensitive to the variation of merging thresholds*. This is thanks to the interaction between merging and the subsequent grouping and re-partitioning. Specifically, when the merging threshold increases, the number of fragments involved in merging decreases, limiting the fragments throughput consequently. This enables grouping and re-partitioning to unlock the potential of re-alignment, i.e., grouping more fragments and building larger batches. Otherwise, as the threshold decreases, the generated fragments are facilitated with larger batch sizes as well as higher resource efficiency, which, however, constrains the performance of grouping and re-partitioning due to SLO requirements.

Notably, based on our empirical experience, *setting merging thresholds less than DNNs' resource margin helps Graft maintain the trade-off between time efficiency and resource efficiency*. The reason is that this lower value compels Graft to sufficiently merge uniform fragments while retaining reasonable optimization space for grouping and re-partitioning. This is in contrast to Uniform, which neglects preservation. Take Res, which has a resource margin of 0.3, an instance. As shown in Figure 16a, Uniform⁺ facilitates Graft with resource reduction by 2.9 \times when its merging threshold decreases from 0.4 to 0.2. Meanwhile, it decreases the number of fragments by 14.5% on average for the subsequent two steps. However, as depicted in Figure 16b, an excessive decrease in the threshold may incur considerable time costs during merging due to its incremental exploration (detailed in §4.1). Empirically, we set the merging threshold at 0.2 by default considering its satisfactory performance in striking the aforementioned trade-off.

5.6 Effectiveness of Grouping

We compare the resource consumption of Graft when adopting the proposed similarity-based grouping and the optimal grouping, respectively, to provision 25 fragments. The




Fig. 16: (a) Resource consumption under varying merging thresholds and varying numbers of DNN fragments, normalized by that under merging threshold of 0.1. (The vertical dashed line is the resource margin, which for VGG and ViT is 0.58 and 3, respectively.) (b) Time costs of Uniform⁺ merging when operating 25 fragments of Res under varying thresholds.




Fig. 17: (a) Resource consumption and time cost of Graft provisioning Inc under varying group sizes; (b) Resource consumption adopting equal grouping weight normalized by the optimal weight.

proposed grouping utilizes the Euclidean distance as the edge weight to build the graph, instead of the exact amount of resources required by Optimal. The result shows that the *similarity-based grouping achieves comparable performance as more than 0.7% compared with Optimal*. Additionally, the similarity-based grouping ensures low overhead for Graft as detailed in §5.9.

We also analyze the impact of the group size and factor weights (namely for partition point, time budget, and throughput as detailed in §4.2). Taking Inc as an example, we observe from Figure 17a that *increased group sizes make a marginal improvement in resource consumption (top), but bring higher time cost (bottom)*. Empirically, a group size of five is a good trade-off point between time complexity and resource savings.

Figure 17b compares the resource consumption of Graft under equal factor weights and optimal factor weights. We can see that the gap is no more than 4.1% on average. This means the *optimization space for factor weights is quite limited*. Based on our empirical experience, slightly increasing the weight for time budget for models with limited resource margins such as Res leads to higher resource efficiency, due to the smaller diversity of the fragments within each group.

5.7 Achievable Throughput

We compare the maximum achievable throughput of Graft, GSLICE, GSLICE⁺, and Static under restricted resource offerings. In this experiment, we gradually increase the number of fragments until their required resources exceed an upper bound. Figure 18 shows that *Graft achieves higher throughput as 1.45 \times , 1.18 \times , 1.29 \times , 1.65 \times , and 2.38 \times on average for Inc, Res, VGG, Mob, and ViT respectively when compared*

with GSLICE⁺. This is because GSLICE⁺ only conducts optimization (batching) among uniform DNN fragments, and ignores that across heterogeneous fragments, thus receiving limited improvement. On the other hand, *Graft achieves higher throughput as 1.57 \times , 1.29 \times , 1.4 \times , 1.76 \times , and 3.3 \times on average for Inc, Res, VGG, Mob, and ViT, respectively than GSLICE*. The improvement is as 1.57 \times , 1.29 \times , 1.09 \times , 1.79 \times , and 3.32 \times , respectively, as compared with Static.

5.8 Massive-Scale Simulations

We further evaluate the resource consumption of Graft, GSLICE, GSLICE⁺, and Static with simulations. We focus on large-scale scenarios with thousands of fragments, which cannot be handled by our testbed. For high time efficiency, we set the merging threshold as 0.01 for Graft. Figure 19 shows that Graft outperforms the baselines. Specifically, *GSLICE and Static have approximate resource consumption, which 1.8 \times , 1.44 \times , 1.76 \times , 2 \times , and 4 \times on average for Inc, Res, VGG, Mob, and ViT, when compared with Graft*. Meanwhile, *Graft reduces the resource consumption by 18.1%, 10.6%, 13.6%, 34.1%, and 16.1% on average, compared with GSLICE⁺, for the five models, respectively*. This means that Graft still outperforms the pure-batching based allocation (GSLICE⁺) at scale. Note that the low merging threshold incurs excessive merging for DNNs of low resource margins or of low heterogeneity (e.g., Res and ViT), compromising Graft's resource efficiency (discussed in §5.5).

5.9 System Overhead

We evaluate the time cost and memory footprint of Graft and Optimal respectively when re-aligning fragments ranging from ten to fifty. Due to limited space, we only illustrate the time costs of Graft. The time cost of Graft achieves an average reduction by up to 99.6% across all five models, in comparison with Optimal. The reason is that Optimal enumerates all the feasible groupings (e.g., 252 for Inc with ten fragments) to figure out the minimum resource consumption, incurring significant overheads. In contrast, Graft only explores one grouping heuristically and thus guarantees time efficiency while achieving close-to-optimum resource efficiency.

On the other hand, as shown in Figure 20a, *the time cost of Graft increases as the number of fragments grows*. This is because under given group sizes, a higher number of fragment yields more groups. While each grouped fragment can be re-aligned independently, more groups inevitably incur higher time costs. In particular, ViT experiences significant high




Fig. 18: Achievable throughput of Graft, GSLICE, GSLICE+, and Static under varying amounts of resources.




Fig. 19: Resource consumption of Graft, GSLICE, GSLICE+, and Static under large numbers of DNN fragments.




Fig. 20: (a) Time cost of Graft when realigning fragments ranging from ten to fifty; (b) Time costs of Graft when realigning fifty fragments of ViT with pool sizes ranging from one to six.

time costs. This is attributed to its highly heterogeneous time budgets, which hinder the effectiveness of merging and result in more groups. To fix this issue, we can introduce a process pool to realign the grouped fragments concurrently. Figure 20b illustrates the time costs when realigning fifty fragments of ViT, which generate six groups after merging. We can see that the time costs decrease by 1.9 \times as the pool size increases from one to two. Notably, the increase in the pool size has diminishing returns. To guarantee time efficiency, we set the pool size as two by default.

In addition, we also evaluate the memory footprint of Graft during the above process. *The average memory consumption is negligible as 35.69MB, 32.14MB, 29.9MB, 34.4MB, and 25.3MB for Inc, Res, VGG, Mob, and ViT respectively. Moreover, this footprint increases no more than 0.63MB with the number of fragments increasing from 10 to 100.*

5.10 Impact of SLOs

We evaluate the impact of SLOs on the performance of Graft. In our work, SLOs express DNN models’ requirements in end-to-end latency. For ease of exposition, the SLO for each model is specified as a ratio to its individual mobile inference latency and we call this ratio as SLO ratio, which can span from 0.5 to 0.9. Moreover, the more strict performance requirements DNNs have in latency, the lower SLO ratios they should be set. We take an e-commerce application as an




Fig. 21: Resource consumption of Graft, normalized by Optimal, under varying SLO ratios

example. This application integrates multiple DNN models, each with specific performance requirements. It is worth noting that *these requirements vary depending on the business domains* [65]. Specifically, for tasks deemed mission-critical, such as facial recognition for payment purposes, it is imperative to configure the corresponding DNN with relatively lower SLO ratios, to meet its stringent latency requirements. In comparison, tasks related to facial beautification, which exhibit a greater degree of tolerance to latency, can be set with higher SLO ratios.

Here, for the generality, we completely evaluate the impacts of SLO ratio over a broad spectrum, spanning from 0.5 to 0.9. Note that hybrid DL (i.e., Neurosurgeon) may fail to find a feasible partition point when the SLO ratio is too small, e.g., less than 0.7 for Inc. Figure 21 shows that *Graft is insensitive to the change of SLO (ratio), constantly achieving performance approximately as good as Optimal.*

5.11 Performance in energy

We evaluate the energy consumption of Graft, GSLICE, GSLICE+, Static, and Static+ under homogeneous small-scale and large-scale scenarios (with the same settings as in Section 5.2 and Section 5.3), respectively. As depicted in Figure 22a, *thanks to its higher resource efficiency, Graft achieves energy reduction by up to 70.1% when compared with GSLICE and Static in the small-scale case.* This improvement becomes more significant in the large-scale scenario due to the incorporation of a greater number of fragments in realignment. Note that GSLICE+ exhibits lower energy consumption than Graft by 42.7%, 32.2%, and 30.1% on average for Inc, VGG, and ViT, respectively. This is because

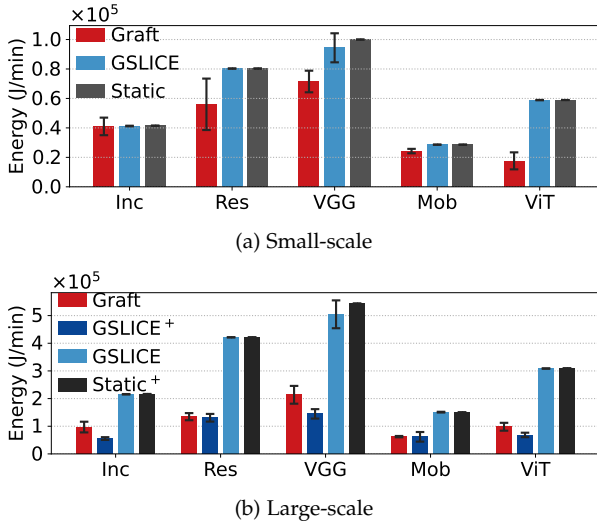


Fig. 22: Energy consumption of Graft, GSLICE, GSLICE⁺, Static, and Static⁺ when realigning (a) four fragments and (b) twenty fragments.

of its sufficient merging enabled by Uniform (detailed in §5.5), which enables GSLICE⁺ with the opportunity to build larger batch sizes and thus reaps higher energy efficiency.

6 LIMITATION AND FUTURE WORK

Realignment disruption. Notably, there may exist cases where the realignment is disrupted by the DNN fragments that undergo frequent variations. As a solution, we could introduce a shadow instances-based alternative as done in GSLICE in the future. Specifically, this strategy sets up shadow instances for the latest arrived DNN fragments when the scheduler is busy with re-aligning existing fragments. Upon completion, the scheduler identifies “similar” fragments, which share the same partition points and approximate time budgets with the recently arrived ones, and then reuses their realignment on the latter. This approach is motivated by the fact that smaller time budgets and higher throughput do not always dictate higher resource consumption due to the discreteness (e.g., batch size) in resource provisioning. Meanwhile, partition points usually occur at several specific layers (illustrated in Figure 7a), presenting a high potential for realignment reuse. The above improvement could allow Graft to effectively handle the frequent variations of DNN fragments.

Intermediate representations with diverse precision. In most cases, Graft does not require specific improvements to deal with intermediate representations’ differences in precision. This is because while intermediate representations may employ low bit-widths for lower transmission latency, most of them can be converted to full precision upon reaching servers [11], thus having no impacts on the subsequent realignment. On the other hand, in cases where the conversion fails to proceed as expected, we can implement the following approach. That is to separate DNN fragments based on their precision while restricting the realignment among the fragments with the same precision.

Availability to other models. Recently early-exiting models are proposed to realize the trade-off between inference accuracy and inference latency [9], [66], [67]. These models

incorporate multiple exits, allowing requests to terminate at different stages of processing. Moreover, prior to execution, requests lack precise information about their exact exiting layer. In particular, when requests exit before reaching the re-partitioning point, the re-aligned DNN fragment will fail to construct predefined batch sizes and thus incur resource over-allocation. To address this issue, we can monitor the achievable throughput (e.g., requests per second) of the exiting layers that are located after the re-partitioning point. Then, Graft adjusts batch sizes accordingly to reap higher resource efficiency.

BERT, LSTM, RNN, and Transformer are designed to analyze sequential data. Compared with CNN, these models have higher data dependency due to their specific requirements of “memory”. This consequently needs to facilitate Graft with an adequate data compression technique to control data exchanging overheads. Notably, to the best of our knowledge, BERT, LSTM, RNN, and Transformer are rarely used alone but rather serve as basic building blocks to construct more complex architectures. Moreover, for the sake of higher accuracy, modern models tend to integrate considerable numbers of blocks, for example, ViT consists of twelve Transformer based blocks. This offers the opportunity to focus on the realignment of inter-blocks, rather than intra-blocks, so as to strike the trade-off between resource efficiency and time complexity. As shown in Table 3, Graft reaps resource reduction by up to 70% for ViT as compared with GSLICE/GSLICE⁺, which only uses batching without realignment.

Heterogeneous models. At present, the realignment is exclusively applied to fragments that originate from the same DNN or, at most, only have slight difference in final layers due to transfer learning. We leave more fine-grained per-layer re-alignment as future work. To support complex scenarios that involves multiple different DNNs, we can introduce a straightforward separation strategy for fragments depending on their DNN types, and then apply Graft.

Split training. Graft demonstrates its effectiveness in provisioning inference tasks on the server side. To support split training, which incorporates both clients and servers to train models collaboratively [30], several considerations must be addressed due to their distinct characteristics compared with inference tasks: (1) Looser time budgets, (2) Higher GPU memory requirements [68], [69], and (3) Larger transmission loads, which encompass batched intermediate data and generated gradients across clients and servers [70]. As a result, Graft needs to adjust its objective. Instead of focusing on minimizing GPU computing resources without violating SLOs, Graft should now prioritize minimizing GPU memory consumption while shortening training duration as much as possible. To achieve this goal, the profiler should collect the profiling data with respect to training throughput and latency in batch sizes and memory consumption. Simultaneously, it is necessary to integrate a series of transmission/storage-targeted optimizations, such as quantization and compression [9], [71], for reaping lower data transmission latency.

Compatible with other hardware. Graft effectively enables fine-grained resource allocation on hardware that supports spatial sharing, e.g., CPUs and GPUs. To seamlessly support temporal-sharing accelerators, Graft requires upgrading its

executor, ensuring that re-aligned DNN fragments are executed in order.

Distributed edge setups. Graft exploits DNN fragments re-alignment within one edge node. To extend to distributed edge setups, it needs to facilitate a straightforward deployment strategy—first bin-packing-based deployment, which is widely adopted by commercial platforms like AWS Lambda and Alibaba Compute Function [72], [73]. Specifically, this strategy selects the first available server that can accommodate the requirements of both the separate and the re-aligned fragments (shown in Figure 4) appropriately. Consequently, it effectively minimizes time-consuming data transmission across servers while mitigating the risk of overwhelming server resources.

7 RELATED WORK

We summarize related work in hybrid DL, DL inference serving, and GPU utilization improvement in this section.

Hybrid DL. Recently, much research has been done on hybrid DL by both academia and industry. Neurosurgeon [23] first introduces DNN partitioning to decrease the end-to-end latency by exploiting the computation capability of a mobile device and cloud servers under dynamic network bandwidth. JointDNN [26] extends the idea to multiple-part partitioning and explores the properties of partition points for different DNN types. Neurosurgeon and JointDNN both focus on DNNs with linear architectures (e.g., AlexNet [74] and VGG [75]). DADS [51] explores partitioning for DNNs based on directed acyclic graph (DAG) (e.g., GoogleNet [76]) and proposes two DNN partitioning strategies for light and heavy network conditions, respectively. There are also studies on specific problems involving hybrid DL such as deploying DNNs at the edge [77], [78], [79], mobile handover [50], and privacy protection [29]. Early-exit has also been merged into DNN partitioning recently for further decreasing the inference latency and resource consumption [9], [66], [67]. Auto-split has been deployed in a commercial cloud to enable edge-cloud collaborative DNN inference [18].

DL inference serving. A large amount of work has been done to improve DL inference serving, aiming at resource efficiency and latency SLO guarantee. The industry develops systems, such as TensorFlow Serving [56] and TorchServe [80]. Meanwhile, the academic proposes a variety of inference serving systems like Clipper [7], Nexus [8], AsyFunc [81], iGniter [82], Opara [83], etc. Clockwork exploits performance predictability and proposes a principled bottom-up approach to achieve low tail latency for inference serving [10]. Systems like ALERT [57] and INFaaS [6] employ model adaptation to balance performance, accuracy, and energy efficiency at runtime. Apart from single models, more complex model graphs have been considered [4], [48], [84]. Different from the above coarse-grain allocation frameworks working at the whole GPU level, GSLICE [59] spatially shares GPUs among located DNNs and hence reaps high throughput with latency guarantee. Besides single DNN inference, there are also some studies on the inference serving “chains” or “graphs” that coordinate multiple DNNs, including GrandSLAM [48] and InferLine [4].

Powerchief [85] focuses on mitigating the latency of multi-stage applications on power-constrained CMP by means of identifying and accelerating the bottleneck service while considering the power budget.

GPU utilization improvement. Numerous spatial sharing based approaches have been proposed to improve GPU utilization. Prophet [86] depends on interference models with respect to shared resources, to realize the “safe” co-location for latency-sensitive applications without violating SLO. Yeung et al. [87] propose a prediction engine to foresee the GPU utilization of heterogeneous DL workloads, to increase resource utilization while reducing performance interference between workloads. FGPU [88] introduces both compute and memory bandwidth isolation mechanisms to promise the predictability of the concurrent applications on GPUs. Slate [89] conducts a workload-aware scheduling for the concurrent kernels to improve the GPU utilization while minimizing resource contention. Maestro [90] explores the GPU resource partitioning at the granularity of both the streaming multi-process (SM) and the simultaneous multi-kernel to improve the system performance. Apart from spatial sharing, temporal sharing and spatial-temporal sharing are also explored to improve GPU utilization [68], [91], [92], [93].

In summary, current work in hybrid DL, such as Neurosurgeon and JointDNN, focuses on partitioning a single DNN to optimize latency or power across one mobile device and one (edge) server, without considering the efficient serving of non-uniform DNN fragments on the server. Meanwhile, existing DL inference serving systems only pay attention to full-size DNNs in the cloud, ignoring the resource efficiency issue in hybrid DL where DNNs are fragmented. In contrast, Graft is optimized for hybrid DL and addresses the inefficiency in provisioning non-uniform fragments, based on a new concept called re-alignment. Graft is orthogonal to GPU spatial- and temporal-sharing improvements, which can be further leveraged to improve the performance of Graft.

8 CONCLUSION

In this paper, we identified the DNN fragment misalignment problem in inference serving for hybrid DL and proposed a new concept called re-alignment to address it by promoting request batching and sharing. We proposed Graft—a first-of-its-kind inference serving system for hybrid DL adopting re-alignment. We identified the challenges in building Graft and proposed efficient algorithms for fine-grained resource allocation for Graft. Experiments based on a system prototype show that Graft achieves significant resource savings while providing latency SLO guarantee.

9 APPENDIX

Bandwidth trace replay. We use the Linux `tc` utility to replay real-world 5G bandwidth trace between mobile devices and the server. Specifically, `tc` offers various queueing disciplines, abbr. `qdisc`, on clients’ network interfaces to shape their traffic. We select Hierarchy Token Bucket-based queueing discipline, i.e., `HTB qdisc`, considering its reliable capacity to ensure specific upper limits on bandwidth [94].

To replay varying 5G network conditions, we encapsulate the control in a script, and execute it periodically (each minute by default as [95]) on each mobile device, which regulates the bandwidth of devices' network interface to predefined values.


Other DNN partition strategies. Due to simplicity, we leverage Neurosurgeon to construct the hybrid DLs context. Specifically, Neurosurgeon concentrates on fundamental factors, namely network conditions, SLO requirements, and mobile device computing capacity, to decide one particular partition point. Other influential factors, such as performance interference across simultaneous workloads [96] and DVFS settings [97] on mobile devices are also investigated. On the other hand, there are works that enable multiple partition points [26], [98]. This offers the potential to achieve lower end-to-end latency but at the cost of higher resource consumption on the server-side [11]. The reason is that multiple partition points dictate extra transmission latency, which compresses the time budgets for inference tasks on the server side. As compensation, more resources are required for SLO guarantees. This contradicts the goal of Graft, i.e., resource-efficient serving on server sides. Consequently, we focus on one partition point-based hybrid DL. Additionally, finer-grained partitioning for DNNs with complex structures, such as parallel branches, has also been explored [51], [99]. Yet, considering the trade-off between time complexity and resource efficiency, we treat parallel branches as a unit without further partitioning.

REFERENCES

- [1] P. Guo and W. Hu, "Potluck: Cross-application approximate deduplication for computation-intensive mobile applications," in *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2018.
- [2] Z. Liao, Z. Luo, Q. Huang, L. Zhang, F. Wu, Q. Zhang, and Y. Wang, "Smart: Screen-based gesture recognition on commodity mobile devices," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 283–295.
- [3] X. Liu, D. Liu, J. Zhang, T. Gu, and K. Li, "Rfid and camera fusion for recognition of human-object interactions," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '21. Association for Computing Machinery, 2021, p. 296–308.
- [4] D. Crankshaw, G. Sela, C. Zumar, X. Mo, J. E. Gonzalez, I. Stoica, and A. Tumanov, "Inferline: latency-aware provisioning and scaling for prediction serving pipelines," in *Proceedings of the 11th ACM Symposium on Cloud Computing*, 2020.
- [5] T. Ben-Nun and T. Hoefler, "Demystifying parallel and distributed deep learning: An in-depth concurrency analysis," *ACM Computing Surveys*, vol. 52, no. 4, pp. 65:1–65:43, 2019.
- [6] F. Romero, Q. Li, N. J. Yadwadkar, and C. Kozyrakis, "INFaaS: Automated model-less inference serving," in *Proceedings of USENIX Annual Technical Conference*, 2021.
- [7] D. Crankshaw, X. Wang, G. Zhou, M. J. Franklin, J. E. Gonzalez, and I. Stoica, "Clipper: A low-latency online prediction serving system," in *Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation*, 2017.
- [8] H. Shen, L. Chen, Y. Jin, L. Zhao, B. Kong, M. Philipose, A. Krishnamurthy, and R. Sundaram, "Nexus: A GPU cluster engine for accelerating DNN-based video analysis," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019.
- [9] S. Laskaridis, S. I. Venieris, M. Almeida, I. Leontiadis, and N. D. Lane, "SPINN: synergistic progressive inference of neural networks over device and cloud," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020.
- [10] A. Gujarati, R. Karimi, S. Alzayat, W. Hao, A. Kaufmann, Y. Vigfusson, and J. Mace, "Serving dnns like clockwork: Performance predictability from the bottom up," in *Proceedings of the 14th USENIX Symposium on Operating Systems Design and Implementation*, 2020.
- [11] M. Almeida, S. Laskaridis, S. I. Venieris, I. Leontiadis, and N. D. Lane, "Dyno: Dynamic onloading of deep neural networks from cloud to device," *ACM Transactions on Embedded Computing Systems*, vol. 21, no. 6, oct 2022. [Online]. Available: <https://doi.org/10.1145/3510831>
- [12] M. Almeida, S. Laskaridis, A. Mehrotra, L. Dudziak, I. Leontiadis, and N. D. Lane, "Smart at what cost? characterising mobile deep neural networks in the wild," ser. IMC '21. New York, NY, USA: Association for Computing Machinery, 2021.
- [13] V. J. Reddi, D. Kanter, P. Mattson, J. Duke, T. Nguyen, R. Chukka, K. Shiring, K. Tan, M. Charlebois, W. Chou, M. El-Khomy, J. Hong, T. S. John, C. Trinh, M. Buch, M. Mazumder, R. Markovic, T. Attafosu, F. Çakir, M. Charkhabi, X. Chen, C. Chiang, D. Dexter, T. Heo, G. Schmuelling, M. Shabani, and D. Zika, "Mlperf mobile inference benchmark: An industry-standard open-source machine learning benchmark for on-device AI," in *Proceedings of Machine Learning and Systems 2022, MLSys 2022, Santa Clara, CA, USA, August 29 - September 1, 2022*, D. Marculescu, Y. Chi, and C. Wu, Eds. mlsys.org, 2022, pp. 352–369.
- [14] A. Ignatov, R. Timofte, A. Kulik, S. Yang, K. Wang, F. Baum, M. Wu, L. Xu, and L. V. Gool, "AI benchmark: All about deep learning on smartphones in 2019," in *2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27–28, 2019*.
- [15] H. Cai, J. Lin, Y. Lin, Z. Liu, H. Tang, H. Wang, L. Zhu, and S. Han, "Enable deep learning on mobile devices: Methods, systems, and applications," *ACM Transactions on Design Automation of Electronic Systems*, 2022.
- [16] Y. Gong, Z. Jiang, Y. Feng, B. Hu, K. Zhao, Q. Liu, and W. Ou, "Edgerec: Recommender system on edge in mobile taobao," *International Conference on Information and Knowledge Management, Proceedings*, pp. 2477–2484, 2020.
- [17] Z. Chen, J. Yao, F. Wang, K. Jia, B. Han, W. Zhang, and H. Yang, "Mc²-sf: Slow-fast learning for mobile-cloud collaborative recommendation," *CoRR*, vol. abs/2109.12314, 2021.
- [18] A. Banitalebi-Dehkordi, N. Vedula, J. Pei, F. Xia, L. Wang, and Y. Zhang, "Auto-split: A general framework of collaborative edge-cloud ai," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery; Data Mining*, ser. KDD '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 2543–2553. [Online]. Available: <https://doi.org/10.1145/3447548.3467078>
- [19] <https://github.com/pytorch/glow/blob/master/docs/Partitioner.md>, 2022.
- [20] https://awsdocs-neuron.readthedocs-hosted.com/en/latest/src/examples/pytorch/resnet50_partition.html, 2022.
- [21] https://sagemaker.readthedocs.io/en/stable/api/training/smd_model_parallel.html, 2022.
- [22] <https://github.com/msr-fiddle/dnn-partitioning>, 2022.
- [23] Y. Kang, J. Hauswald, C. Gao, A. Rovinski, T. N. Mudge, J. Mars, and L. Tang, "Neurosurgeon: Collaborative intelligence between the cloud and mobile edge," in *Proceedings of the 22nd ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2017.
- [24] J. Huang, C. Samplawski, D. Ganesan, B. M. Marlin, and H. Kwon, "CLIO: enabling automatic compilation of deep learning pipelines across iot and cloud," in *MobiCom '20: The 26th Annual International Conference on Mobile Computing and Networking, London, United Kingdom, September 21–25, 2020*. ACM, 2020, pp. 58:1–58:12.
- [25] L. Zhang, L. Chen, and J. Xu, "Autodidactic neurosurgeon: Collaborative deep inference for mobile edge intelligence via online learning," ser. WWW '21. New York, NY, USA: Association for Computing Machinery, 2021.
- [26] A. E. Eshratifar, M. S. Abrishami, and M. Pedram, "Jointdnn: An efficient training and inference engine for intelligent mobile cloud computing services," *IEEE Transactions on Mobile Computing*, vol. 20, no. 2, pp. 565–576, 2021.
- [27] J. Wu, L. Wang, Q. Pei, X. Cui, F. Liu, and T. Yang, "Hitdl: High-throughput deep learning inference at the hybrid mobile edge," *IEEE Transactions Parallel Distributed System*, vol. 33, no. 10, pp. 4499–4514, 2022.
- [28] S. Zhang, Y. Li, X. Liu, S. Guo, W. Wang, J. Wang, B. Ding, and D. Wu, "Towards real-time cooperative deep inference over the

- cloud and edge end devices," *Proceedings ACM Interactive Mobile Wearable and Ubiquitous Technologies*, vol. 4, no. 2, pp. 69:1–69:24, 2020. [Online]. Available: <https://doi.org/10.1145/3397315>
- [29] F. Mireshghallah, M. Taram, P. Ramrakhiani, A. Jalali, D. M. Tullsen, and H. Esmaeilzadeh, "Shredder: Learning noise distributions to protect inference privacy," in *ASPLOS '20: Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, March 16-20, 2020*, J. R. Larus, L. Ceze, and K. Strauss, Eds. ACM, 2020, pp. 3–18.
- [30] M. G. Poirot, P. Vepakomma, K. Chang, J. Kalpathy-Cramer, R. Gupta, and R. Raskar, "Split learning for collaborative deep learning in healthcare," *CoRR*, vol. abs/1912.12115, 2019.
- [31] L. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017.
- [32] Y. Li, A. Phanishayee, D. Murray, and N. S. Kim, "Doing more with less: training large DNN models on commodity servers for the masses," in *Proceedings of the 18th Workshop on Hot Topics in Operating Systems*, 2021.
- [33] P. with code, "Gflops of image classification targeted models on imagenet," <https://paperswithcode.com/sota/image-classification-on-imagenet?metric=GFLOPs>, 2022.
- [34] W. Dong, J. Lv, G. Chen, Y. Wang, H. Li, Y. Gao, and D. Bharadia, "TinyNet: A lightweight, modular, and unified network architecture for the internet of things," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, ser. MobiSys '22, 2022, p. 248–260.
- [35] S. Jiang, Z. Lin, Y. Li, Y. Shu, and Y. Liu, "Flexible high-resolution object detection on edge devices with tunable latency," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '21, 2021, p. 559–572.
- [36] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, "Once-for-all: Train one network and specialize it for efficient deployment," in *Proceedings of the 8th International Conference on Learning Representations*, 2020, pp. 1–11.
- [37] T. Zhang, S. Ye, K. Zhang, J. Tang, W. Wen, M. Fardad, and Y. Wang, "A systematic DNN weight pruning framework using alternating direction method of multipliers," in *Proceedings of the 15th European Conference on Computer Vision*, 2018.
- [38] P. Nayak, D. Zhang, and S. Chai, "Bit efficient quantization for deep neural networks," in *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing - NeurIPS Edition*, 2019.
- [39] A. Potapov and S. Rodionov, "Genetic algorithms with dnn-based trainable crossover as an example of partial specialization of general search," in *Proceedings of the 10th International Conference on Artificial General Intelligence*, 2017.
- [40] F. Jia, D. Zhang, T. Cao, S. Jiang, Y. Liu, J. Ren, and Y. Zhang, "Codl: Efficient cpu-gpu co-execution for deep learning inference on mobile devices," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, ser. MobiSys '22, 2022, p. 209–221.
- [41] J. Park, K. Bin, and K. Lee, "Mgemm: Low-latency convolution with minimal memory overhead optimized for mobile devices," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, 2022, p. 222–234.
- [42] M. Wang, S. Ding, T. Cao, Y. Liu, and F. Xu, "Asymo: Scalable and efficient deep-learning inference on asymmetric mobile cpus," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. Association for Computing Machinery, 2021, p. 215–228.
- [43] R. Han, Q. Zhang, C. H. Liu, G. Wang, J. Tang, and L. Y. Chen, "Legodnn: Block-grained scaling of deep neural networks for mobile vision," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '21, 2021, p. 406–419.
- [44] M. Xu, J. Liu, Y. Liu, F. X. Lin, Y. Liu, and X. Liu, "A first look at deep learning apps on smartphones," in *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, L. Liu, R. W. White, A. Mantrach, F. Silvestri, J. J. McAuley, R. Baeza-Yates, and L. Zia, Eds. ACM, 2019, pp. 2125–2136. [Online]. Available: <https://doi.org/10.1145/3308558.3313591>
- [45] B. Lu, J. Yang, W. Jiang, Y. Shi, and S. Ren, "One proxy device is enough for hardware-aware neural architecture search," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 5, no. 3, pp. 34:1–34:34, 2021. [Online]. Available: <https://doi.org/10.1145/3491046>
- [46] C. Wu, D. Brooks, K. Chen, D. Chen, S. Choudhury, M. Dukhan, K. M. Hazelwood, E. Isaac, Y. Jia, B. Jia, T. Leyvand, H. Lu, Y. Lu, L. Qiao, B. Reagen, J. Spisak, F. Sun, A. Tulloch, P. Vajda, X. Wang, Y. Wang, B. Wasti, Y. Wu, R. Xian, S. Yoo, and P. Zhang, "Machine learning at facebook: Understanding inference at the edge," in *25th IEEE International Symposium on High Performance Computer Architecture, HPCA 2019, Washington, DC, USA, February 16-20, 2019*. IEEE, 2019, pp. 331–344.
- [47] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017. [Online]. Available: <https://doi.org/10.1109/MC.2017.9>
- [48] R. S. Kannan, L. Subramanian, A. Raju, J. Ahn, J. Mars, and L. Tang, "Grandslam: Guaranteeing slas for jobs in microservices execution frameworks," in *Proceedings of the 14th EuroSys Conference*, 2019.
- [49] T. Shen, J. Qi, J. Jiang, X. Wang, X. Wen, X. Chen, and S. Zhao, "Soter: Guarding black-box inference for general neural networks at the edge," in *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. Carlsbad, CA: USENIX Association, Jul. 2022. [Online]. Available: <https://www.usenix.org/conference/atc22/presentation/shen>
- [50] W. Ju, D. Yuan, W. Bao, L. Ge, and B. B. Zhou, "Deepsave: Saving dnn inference during handovers on the edge," in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 2019.
- [51] C. Hu, W. Bao, D. Wang, and F. Liu, "Dynamic adaptive DNN surgery for inference acceleration on the edge," in *Proceedings of the 38th IEEE International Conference on Computer Communications*, 2019.
- [52] J. Emmons, S. Fouladi, G. Ananthanarayanan, S. Venkataraman, S. Savarese, and K. Winstein, "Cracking open the dnn black-box: Video analytics with dnns across the camera-cloud boundary," ser. HotEdgeVideo'19, 2019.
- [53] <https://www.huaweicloud.com/intl/en-us/solution/edgecloud.html>, 2022.
- [54] Y. Huang, X. Qiao, S. Dustdar, and Y. Li, "Aodnn: An auto-offloading approach to optimize deep inference for fostering mobile web," in *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, 2022.
- [55] D. Raca, D. Leahy, C. J. Sreenan, and J. J. Quinlan, "Beyond throughput, the next generation: a 5g dataset with channel and context metrics," in *Proceedings of the 11th ACM Multimedia Systems Conference*, 2020.
- [56] <https://www.tensorflow.org/tfx/guide/serving#introduction>, 2020.
- [57] C. Wan, M. H. Santriaji, E. Rogers, H. Hoffmann, M. Maire, and S. Lu, "ALERT: accurate learning for energy and timeliness," in *Proceedings of USENIX Annual Technical Conference*, 2020.
- [58] J. Li, L. Zhao, Y. Yang, K. Zhan, and K. Li, "Tetris: Memory-efficient serverless inference through tensor sharing," in *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. Carlsbad, CA: USENIX Association, Jul. 2022. [Online]. Available: <https://www.usenix.org/conference/atc22/presentation/li-jie>
- [59] A. Dhakal, S. G. Kulkarni, and K. Ramakrishnan, "Gslice: controlled spatial sharing of gpus for a scalable inference platform," in *Proceedings of the 11th ACM Symposium on Cloud Computing*, 2020.
- [60] K. Andreev and H. Racke, "Balanced graph partitioning," *Theory of Computing Systems*, vol. 39, no. 6, pp. 929–939, 2006.
- [61] C. E. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic, "FENNEL: streaming graph partitioning for massive scale graphs," in *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*, 2014. ACM, 2014, pp. 333–342. [Online]. Available: <https://doi.org/10.1145/2556195.2556213>
- [62] <https://www.gurobi.com/resource/linear-programming-basics/>, 2022.
- [63] <https://piverval.com/jetson-nano-power-supply-guide/>, 2023.
- [64] https://developer.ridgerun.com/wiki/index.php/NVIDIA_Jetson_TX2_NVP_model, 2023.
- [65] S. Luo, H. Xu, K. Ye, G. Xu, L. Zhang, J. He, G. Yang, and C. Xu, "Erms: Efficient resource management for shared microservices with SLA guarantees," in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1, ASPLOS 2023, Vancouver, BC, Canada, March 25-29, 2023*. ACM, 2023, pp. 62–77.
- [66] E. Li, L. Zeng, Z. Zhou, and X. Chen, "Edge AI: on-demand accelerating deep neural network inference via edge computing," *IEEE Transactions on Wireless Communications*, vol. 19, no. 1, pp. 447–457, 2020.

- [67] Y. Li, C. Zhang, S. Han, L. L. Zhang, B. Yin, Y. Liu, and M. Xu, "Boosting mobile CNN inference through semantic memory," in *ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*. ACM, 2021, pp. 2362–2371.
- [68] P. Yu and M. Chowdhury, "Salus: Fine-grained GPU sharing primitives for deep learning applications," *CoRR*, vol. abs/1902.04610, 2019. [Online]. Available: <http://arxiv.org/abs/1902.04610>
- [69] G. Lim, J. Ahn, W. Xiao, Y. Kwon, and M. Jeon, "Zico: Efficient GPU memory sharing for concurrent DNN training," in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, 2021, pp. 161–175.
- [70] Z. Dong, Q. He, F. Chen, H. Jin, T. Gu, and Y. Yang, "Edgemove: Pipelining device-edge model training for mobile intelligence," in *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*. ACM, 2023, pp. 3142–3153.
- [71] Y. Bai, C. Li, Q. Zhou, J. Yi, P. Gong, F. Yan, R. Chen, and Y. Xu, "Gradient compression supercharged high-performance data parallel DNN training," in *SOSP '21: ACM SIGOPS 28th Symposium on Operating Systems Principles, Virtual Event / Koblenz, Germany, October 26–29, 2021*, R. van Renesse and N. Zeldovich, Eds. ACM, 2021, pp. 359–375.
- [72] H. Tian, S. Li, A. Wang, W. Wang, T. Wu, and H. Yang, "Owl: performance-aware scheduling for resource-efficient function-as-a-service cloud," in *Proceedings of the 13th Symposium on Cloud Computing, SoCC 2022, San Francisco, California, November 7–11, 2022*, A. Gavrilovska, D. Altinbükten, and C. Binnig, Eds. ACM, 2022, pp. 78–93.
- [73] L. Wang, M. Li, Y. Zhang, T. Ristenpart, and M. M. Swift, "Peeking behind the curtains of serverless platforms," in *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11–13, 2018*, H. S. Gunawi and B. C. Reed, Eds. USENIX Association, 2018, pp. 133–146.
- [74] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proceedings of the 26th Annual Conference on Neural Information Processing Systems*, 2012.
- [75] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proceedings of the 3rd International Conference on Learning Representations*, 2015, pp. 1–14.
- [76] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [77] K. Hsu, K. Bhardwaj, and A. Gavrilovska, "Couper: DNN model slicing for visual analytics containers at the edge," in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 2019.
- [78] H. Jeong, H. Lee, C. H. Shin, and S. Moon, "IONN: incremental offloading of neural network computations from mobile devices to edge servers," in *Proceedings of the 9th ACM Symposium on Cloud Computing*, 2018.
- [79] Q. Chen, Z. Zheng, C. Hu, D. Wang, and F. Liu, "On-edge multi-task transfer learning: Model and practice with data-driven task allocation," *IEEE Transactions Parallel Distributed System*, vol. 31, no. 6, pp. 1357–1371, 2020.
- [80] <https://pytorch.org/serve/>, 2020.
- [81] Q. Pei, Y. Yuan, H. Hu, Q. Chen, and F. Liu, "Asyfunc: A high-performance and resource-efficient serverless inference system via asymmetric functions," in *Proceedings of the 2023 ACM Symposium on Cloud Computing, SoCC 2023, Santa Cruz, CA, USA, 30 October 2023 - 1 November 2023*. ACM, 2023, pp. 324–340.
- [82] F. Xu, J. Xu, J. Chen, L. Chen, R. Shang, Z. Zhou, and F. Liu, "igniter: Interference-aware GPU resource provisioning for predictable DNN inference in the cloud," *IEEE Transactions Parallel Distributed System*, vol. 34, no. 3, pp. 812–827, 2023.
- [83] A. Chen, F. Xu, Y. Dong, L. Chen, Z. Zhou, and F. Liu, "Opara: Exploring operator parallelism for expediting dnn inference on gpus," China Computer Federation System, Nan Change, China, Tech. Rep., August 2023. [Online]. Available: <https://github.com/icloud-ecnu/Opara/blob/main/pdf/tc-opara.pdf>
- [84] F. Romero, M. Zhao, N. J. Yadwadkar, and C. Kozyrakis, "Llama: A Heterogeneous & Serverless Framework for Auto-Tuning Video Analytics Pipelines," *arXiv preprint arXiv:2102.01887*, 2021.
- [85] H. Y. and Quan Chen, M. Riaz, Z. Luan, L. Tang, and J. Mars, "Powerchief: Intelligent power allocation for multi-stage applications to improve responsiveness on power constrained CMP," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, 2017.
- [86] Q. Chen, H. Yang, M. Guo, R. S. Kannan, J. Mars, and L. Tang, "Prophet: Precise qos prediction on non-preemptive accelerators to improve utilization in warehouse-scale computers." ACM, 2017, pp. 17–32. [Online]. Available: <https://doi.org/10.1145/3037697.3037700>
- [87] G. Yeung, D. Borowiec, A. Friday, R. Harper, and P. Garraghan, "Towards GPU utilization prediction for cloud deep learning," in *12th USENIX Workshop on Hot Topics in Cloud Computing, HotCloud 2020, July 13–14, 2020*, A. Phanishayee and R. Stutsman, Eds. USENIX Association, 2020. [Online]. Available: <https://www.usenix.org/conference/hotcloud20/presentation/yeung>
- [88] S. Jain, I. Baek, S. Wang, and R. Rajkumar, "Fractional gpus: Software-based compute and memory bandwidth reservation for gpus," in *25th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2019, Montreal, QC, Canada, April 16–18, 2019*. IEEE, 2019, pp. 29–41. [Online]. Available: <https://doi.org/10.1109/RTAS.2019.00011>
- [89] T. Allen, S. Feng, and R. Ge, "Slate: Enabling workload-aware efficient multiprocessing for modern gpgpus," in *2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2019, Rio de Janeiro, Brazil, May 20–24, 2019*. IEEE, 2019, pp. 252–261. [Online]. Available: <https://doi.org/10.1109/IPDPS.2019.00035>
- [90] J. J. K. Park, Y. Park, and S. A. Mahlke, "Dynamic resource management for efficient utilization of multitasking gpus," in *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2017, Xi'an, China, April 8–12, 2017*. ACM, 2017, pp. 527–540. [Online]. Available: <https://doi.org/10.1145/3037697.3037707>
- [91] Y. Hu, S. Rallapalli, B. Ko, and R. Govindan, "Olympian: Scheduling GPU usage in a deep neural network model serving system," in *Proceedings of the 19th International Middleware Conference, Middleware 2018, Rennes, France, December 10–14, 2018*. ACM, 2018, pp. 53–65. [Online]. Available: <https://doi.org/10.1145/3274808.3274813>
- [92] G. Chen, Y. Zhao, X. Shen, and H. Zhou, "Effisha: A software framework for enabling efficient preemptive scheduling of gpu," in *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, ser. PPOPP '17*. New York, NY, USA: Association for Computing Machinery, 2017, p. 3–16. [Online]. Available: <https://doi.org/10.1145/3018743.3018748>
- [93] S. Choi, S. Lee, Y. Kim, J. Park, Y. Kwon, and J. Huh, "Serving heterogeneous machine learning models on multi-gpu servers with spatio-temporal sharing," in *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. Carlsbad, CA: USENIX Association, 2022. [Online]. Available: <https://www.usenix.org/conference/atc22/presentation/choi-seungbeom>
- [94] "Tc classful qdiscs," https://man7.org/linux/man-pages/man8/tc.8.html#CLASSFUL_QDISCS, 2023.
- [95] X. Fu, T. Ghaffar, J. C. Davis, and D. Lee, "Edgewise: A better stream processing engine for the edge," in *2019 USENIX Annual Technical Conference, USENIX ATC 2019, Renton, WA, USA, July 10–12, 2019*. USENIX Association, 2019, pp. 929–946.
- [96] Y. G. Kim and C. Wu, "Autoscale: Energy efficiency optimization for stochastic edge inference using reinforcement learning," in *53rd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2020, Athens, Greece, October 17–21, 2020*. IEEE, 2020, pp. 1082–1096.
- [97] Z. Zhang, Y. Zhao, and J. Liu, "Poster abstract: Dvfo: Dynamic voltage, frequency and offloading for efficient ai on edge devices," in *Proceedings of the 22nd International Conference on Information Processing in Sensor Networks, ser. IPSN '23*, New York, NY, USA, 2023, p. 304–305.
- [98] P. Ren, X. Qiao, Y. Huang, L. L. C. Pu, and S. Dustdar, "Fine-grained elastic partitioning for distributed dnn towards mobile web ar services in the 5g era," *IEEE Transactions on Services Computing*, vol. 15, no. 6, pp. 3260–3274, 2022.
- [99] K. Fu, J. Shi, Q. Chen, N. Zheng, W. Zhang, D. Zeng, and M. Guo, "Qos-aware irregular collaborative inference for improving throughput of DNN services," in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis, Dallas, TX, USA, November 13–18, 2022*, F. Wolf, S. Shende, C. Culhane, S. R. Alam, and H. Jagode, Eds. IEEE, 2022, pp. 69:1–69:14.



Jing Wu received the M.S. degree in the School of Computer Science and Engineering, Northeastern University, Shenyang, China, in 2018. She is currently a Ph.D. student in the School of Computer Science and Technology, Huazhong University of Science and Technology, China. Her research interests include edge computing, augmented reality, and deep learning.



Lin Wang is currently a Full Professor and Head of the Chair of Computer Networks in the Department of Computer Science at Paderborn University. He received his Ph.D. degree from the Institute of Computing Technology, Chinese Academy of Sciences in 2015. Previously, he held positions at Vrije Universiteit Amsterdam, TU Darmstadt, SnT Luxembourg, and IMDEA Networks Institute. He is broadly interested in networked systems, with a focus on in-network computing and intermittently-powered IoT systems.

He has received a Google Research Scholar Award, an Outstanding Paper Award of RTSS 2022, a Best Paper Award of IPCCC 2023, and an Athene Young Investigator Award of TU Darmstadt. He is currently a Senior Member of IEEE.



Qirui Jin is a current undergraduate student at the University of Michigan — Ann Arbor. His research topics include machine learning systems, reverse engineering, and neural network verification.



Fangming Liu (S'08, M'11, SM'16) received the B.Eng. degree from the Tsinghua University, Beijing, and the Ph.D. degree from the Hong Kong University of Science and Technology, Hong Kong. He is currently a Full Professor with the Huazhong University of Science and Technology, Wuhan, China. His research interests include cloud computing and edge computing, data-center and green computing, SDN/NFV/5G and applied ML/AI. He received the National Natural Science Fund (NSFC) for Excellent Young

Scholars, and the National Program Special Support for Top-Notch Young Professionals. He is a recipient of the Best Paper Award of IEEE/ACM IWQoS 2019, ACM e-Energy 2018 and IEEE GLOBECOM 2011, the First Class Prize of Natural Science of Ministry of Education in China, as well as the Second Class Prize of National Natural Science Award in China.