

# Collaborative Processing for Multi-Tenant Inference on Memory-Constrained Edge TPUs

Nathan Ng, Walid A. Hanafy, Prashanthi Kadambi<sup>†</sup>, Balachandra Sunil,  
Ayush Gupta, David Irwin, Yogesh Simmhan<sup>†</sup>, Prashant Shenoy  
University of Massachusetts Amherst  
<sup>†</sup>Indian Institute of Science

**Abstract**—IoT applications increasingly rely on on-device AI accelerators to ensure high performance, especially in low-connectivity and safety-critical scenarios. However, the limited on-chip memory of these accelerators forces inference runtimes to swap model segments between host and accelerator memory, incurring significant swapping overheads. While collaborative processing by partitioning model execution across CPU and accelerator resources can reduce accelerator memory pressure and execution overhead, naive partitioning may worsen end-to-end latency by either shifting excessive computation to the CPU or failing to sufficiently reduce swapping, a problem that is further exacerbated in multi-tenant and dynamic environments.

To address these issues, we present SwapLess, a system for adaptive, multi-tenant TPU-CPU collaborative inference on memory-constrained Edge TPUs. SwapLess utilizes an analytic queuing model that captures partition-dependent CPU/TPU service times as well as inter- and intra-model swapping overheads across different workload mixes and request rates. Using this model, SwapLess continuously adjusts both the partition point and CPU core allocation online to minimize end-to-end response time with low decision overhead. An implementation on Edge TPU-equipped platforms demonstrates that SwapLess reduces mean latency by up to 63.8% for single-tenant workloads and up to 77.4% for multi-tenant workloads relative to the default Edge TPU compiler.

**Index Terms**—Memory-Constrained Accelerators, Model Partitioning, Multi-Tenant Inference, Performance Modeling

## I. INTRODUCTION

IoT applications increasingly rely on AI models to interpret and respond to their environments, such as smart cameras performing real-time activity recognition, wearables inferring health and mobility signals, and industrial sensors detecting anomalies for predictive maintenance [1]. While a common architectural approach is to offload computation from resource-constrained devices to nearby edge servers or the cloud, many IoT deployments cannot rely on offloading due to intermittent network connectivity or limited bandwidth [2], [3]. Furthermore, in safety-critical and tactical scenarios such as disaster response and military operations, communication may be disrupted or restricted for security reasons [4], [5]. In such settings, IoT systems must execute inference locally, commonly referred to as on-device AI [1].

To support the growing demands of AI models, IoT systems often rely on programmable AI accelerators, such as Google’s Edge Tensor Processing Unit (Edge TPU) [6] and Raspberry Pi AI HAT [7], which enable low-end devices to perform

tasks such as image classification and object detection with reasonable performance comparable to edge offloading. Despite these advances, many accelerators have limited on-chip memory. For example, Edge TPUs provide only 8 MB of SRAM for caching model parameters, which is insufficient for storing models such as InceptionV4 or ResNet50, resulting in high inference latency.

When a model exceeds TPU memory capacity, the inference runtime relies on memory-swapping approaches that sequentially swap model segments between host and accelerator memory during execution, introducing significant latency overhead. For instance, our analysis in Fig. 1 shows that swapping overhead can account for up to 62.4% of the observed latency. This problem is further exacerbated in multi-tenant scenarios where multiple models run concurrently on the same device, increasing both the combined memory footprint and swap frequency. To address this issue, prior research has focused on multi-TPU settings where models are split across TPU instances based on model size and request rates [8]–[11].

Instead of scaling out to multiple accelerators, an alternative is to leverage the host CPU cores available on these edge platforms. Prior work has shown that, for the less-parallelizable layers of deep neural networks, general-purpose CPUs can deliver inference performance comparable to specialized accelerators, which often derive their advantage primarily from highly parallel layers [12]. Building on this insight, we propose a collaborative processing approach that selectively offloads the trailing layers to the host CPU, providing a practical lever to reduce memory pressure and swapping overhead on the TPU. However, as we show in Fig. 7, naive partitioning may result in higher end-to-end latency than using the TPU alone. For example, choosing a partition that ensures the TPU-resident portion fits entirely in accelerator memory can substantially increase end-to-end latency by shifting too much computation to the CPU. Conversely, a conservative partition that offloads only the trailing layers may incur excessive swapping overhead.

To minimize end-to-end response time, a partitioning strategy must account for the relative processing speed of each segment on the CPU and TPU, the memory footprint of the TPU-resident segment and its associated swapping overhead, and the request arrival rate. These challenges are further exacerbated in multi-tenant, dynamic settings, where both the set




Fig. 1. Intra-model swapping overhead can contribute up to 62.4% of total TPU inference latency.




Fig. 2. Inter-model swapping overhead can contribute up to 49% of total latency.




Fig. 3. Comparable performance in later segments offers opportunities for collaborative inference.

of active models and their request rates can change over time. For example, the partitioning and request rate of one model directly affect the load on both the TPU and CPU, which in turn influence the queuing delays and swapping overheads experienced by other models. This interdependence requires that offloading decisions and resource assignments be made holistically across all concurrently executing models, making model partitioning and CPU core allocation a complex, multi-dimensional resource allocation problem.

To address these challenges, we propose SwapLess, a system for collaborative processing of multi-tenant inference on memory-constrained Edge TPU devices. SwapLess continuously adjusts partitioning and CPU resource allocation online to minimize end-to-end response time with negligible decision overhead. In designing, implementing, and evaluating SwapLess, we make the following contributions:

- At the core of SwapLess decisions is an analytic queuing model that characterizes the performance of collaborative TPU-CPU inference in multi-tenant settings. The model jointly captures the effects of model partitioning, resource allocation, and inter- and intra-model swapping overheads across workload mixes and request rates.
- We present the design of SwapLess, a system that adaptively optimizes model partitioning and CPU core allocation on Edge TPU-equipped devices. SwapLess uses a greedy hill-climbing resource allocation algorithm that leverages our analytic queuing model to minimize inference latency.
- We implement and evaluate SwapLess against three baselines across a range of workloads. Results show that SwapLess reduces mean latency by up to 63.8% for single-tenant workloads and up to 77.4% for multi-tenant workloads compared to the default Edge TPU compiler.

## II. BACKGROUND AND MOTIVATION

This section provides background on Edge TPU inference, quantifies the memory swapping overhead in single- and multi-tenant scenarios, and discusses the opportunities and challenges of collaborative TPU-CPU inference.

### A. Edge TPU Hardware and Inference Execution

Google’s Edge TPU is an application-specific integrated circuit (ASIC) with a systolic array architecture designed to accelerate AI inference on low-end devices. Each Edge TPU contains 8MB of on-chip SRAM for caching model

parameters. An AI model is represented as a TensorFlow Lite computation graph and compiled with the Edge TPU compiler to generate optimized TPU-executable code. To run models larger than its on-chip memory, the TPU runtime first executes the model segment currently in memory, then sequentially swaps in the next segment from host memory and executes it, repeating until all segments are processed. In multi-tenant settings, the compiler prioritizes memory allocation based on the order in which models are specified in the compiler command, and some models may receive no initial memory allocation. For these models, execution requires first swapping their weights into memory, which incurs additional overhead compared to models whose weights are already resident on the TPU. Since modern Transformer architectures involve dynamic attention, embedding layers, and non-linear operations that are not natively supported [13], this work focuses on optimizing execution for standard convolutional models.

### B. Memory Swapping Overheads

The Edge TPU’s limited memory poses a key performance bottleneck for many widely used model architectures, including InceptionV4, ResNet50V2, and DenseNet201, which exceed this limit. Fig. 1 illustrates this *intra-model swapping* overhead using a micro-benchmark, where each model is partitioned into segments that fit in TPU memory. We profile the execution time of each segment and compare the total to the full model executed on the TPU, including swapping, to isolate the intra-model swapping overhead. The results reveal that intra-model swapping overhead is a dominant factor in total inference latency, ranging from 20.2% for DenseNet201 to 62.4% for larger architectures such as InceptionV4.

This problem is further exacerbated in multi-tenant settings by *inter-model swapping*, which occurs when the combined memory footprint of concurrent models exceeds the memory limit. When requests for different models are executed in succession, parameters of the subsequent model that are not present in SRAM must be loaded from host memory, evicting existing weights and incurring swapping overhead before execution. Fig. 2 illustrates this overhead across different workload mixes. In the workload mix with MobileNetV2 and SqueezeNet, no swapping occurs because their combined size fits within the memory. However, models with larger combined footprints, such as EfficientNet and GPUNet, experience significant performance degradation. For a balanced

50:50 request mix, each request has roughly a 50% chance of following a request for a different model, requiring an inter-model swap since its weights were evicted by the model used by the previous request. In this case, swapping overhead accounts for up to 35% of total latency compared to standalone execution. When the workload is skewed (e.g., a 90:10 mix), this penalty is even more pronounced for the less-frequent model, accounting for up to 49% of its total execution latency.

### C. Opportunities and Challenges of CPU Offloading

Prior work [12] has shown that GPUs and CPUs can achieve comparable performance in later, less parallelizable layers of AI models. Our experiments across a wide range of models confirm a similar opportunity on Edge TPUs. Fig. 3 compares the performance between TPU and CPU across segments of an InceptionV4 model. As shown, the first segment achieves a substantial TPU performance gain, which decreases in subsequent segments, with the last three segments showing comparable performance on TPU and CPU. We observe a similar trend across other models, and the results are omitted due to space constraints. This observation suggests an opportunity to minimize memory-swapping overhead and inference latency by reducing the TPU memory footprint through offloading model layers to the CPU.

Effectively capitalizing on this opportunity, however, requires careful resource allocation. First, the optimal partition point for each model must balance the relative processing speeds of the TPU and CPU for each segment, the memory footprint of the segments, and the interactions among co-located models, accounting for heterogeneous compute demands and varying memory-swapping overheads that can significantly impact overall performance. Second, in multi-tenant settings, the system must jointly consider model partitioning and CPU core allocation. The processing time of offloaded model segments on the CPU depends on both the number of cores allocated and the amount of computation offloaded, which in turn interacts with TPU memory allocation. Finally, runtime dynamics add further complexity, as request rates fluctuate and models may be added or removed over time. Consequently, resource allocation strategies must be adaptive and lightweight, responding to changing workloads without incurring excessive computational overhead.

## III. SWAPLESS DESIGN

To address the above challenges, we propose SwapLess, a system designed to optimize multi-tenant inference on Edge TPU-equipped devices. This section provides an overview of the SwapLess design, introduces our proposed analytic queueing model that captures the effects of model partitioning, resource allocation, and memory swapping, and outlines its algorithm for joint model partitioning and core allocation.

### A. System Overview

Fig. 4 presents the architecture of SwapLess, which consists of an offline phase and an online phase. In the offline

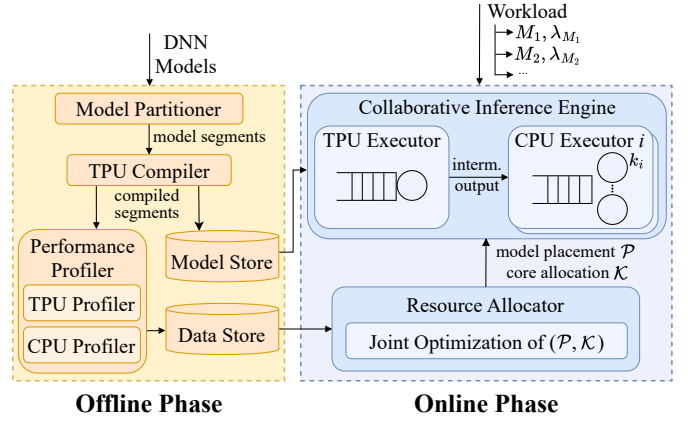


Fig. 4. Overview of SwapLess architecture.

phase, SwapLess analyzes the AI models to identify viable partition points and partitions each model into a set of TPU prefixes and CPU suffixes. It compiles the resulting segments for TPU execution and profiles them on both the TPU and CPU to characterize execution latency and resource usage. The profile results and compiled binaries are stored for use at runtime. During the online phase, the SwapLess resource allocator uses this performance data and an analytic queueing model to jointly determine model partitioning and CPU core allocations, predicting system latency for candidate schemes to optimize end-to-end performance under memory and compute constraints. SwapLess then executes requests according to the optimized allocation, assigning model prefixes to a global TPU worker and forwarding intermediate outputs to dedicated CPU executors, where suffixes are processed using the allocated CPU cores.

### B. Analytic Queueing Model

To enable principled resource allocation, SwapLess accounts for (i) queueing delays at the shared TPU under mixed workloads, (ii) CPU-side contention under per-model core allocations, and (iii) the latency impact of memory swapping. SwapLess uses an analytic model based on queueing theory, a well-established tool for modeling response times in computing systems, to estimate end-to-end latency under a given partitioning and resource allocation scheme. These estimates are then used by SwapLess's resource allocator to make adaptive resource allocation decisions.

Consider a device that concurrently executes a set of AI models  $\mathcal{M} = \{M_1, M_2, \dots, M_n\}$ . Each model  $M_i \in \mathcal{M}$  receives requests according to a Poisson process with arrival rate  $\lambda_{M_i}$ , where  $\Lambda = \{\lambda_{M_1}, \dots, \lambda_{M_n}\}$  denotes the set of all arrival rates. Each model  $M_i$  exposes  $P_i$  candidate partition points that define potential boundaries between TPU and CPU execution. For a selected partition point  $p_i \in \{0, \dots, P_i\}$ , the prefix block  $M_{i,[1:p_i]}$  executes on the Edge TPU, while the suffix block  $M_{i,[p_i+1:P_i]}$  is offloaded to the host CPU. By convention,  $p_i = 0$  corresponds to full CPU execution, and  $p_i = P_i$  corresponds to full TPU execution.

TABLE I  
NOTATIONS USED IN ANALYTIC MODELS.

| Notation                                | Description   |
|---|---|
| <i>Workload and Hardware Parameters</i> |   |
| $\mathcal{M}$                           | Set of $n$ AI models $\{M_1, \dots, M_n\}$ .                            |
| $P_i$                                   | Total candidate partition points for $M_i$ .                            |
| $M_{i,[a:b]}$                           | Block of model $M_i$ from layer $a$ to $b$ .                            |
| $d^{in}, d^{out}$                       | Input and intermediate tensor size.                                     |
| $B$                                     | Memory-to-TPU bandwidth.  |
| $C$                                     | TPU SRAM capacity.  |
| $K_{max}$                               | Total available physical CPU cores.                                     |
| <i>Decision Variables</i>               |   |
| $p_i, k_i$                              | Selected partition point and CPU cores allocated to $M_i$ .             |
| $\mathcal{P}, \mathcal{K}$              | Global model partitioning and core allocation vectors.                  |
| <i>Performance Parameters</i>           |   |
| $\lambda_{M_i}$                         | Request arrival rate for model $M_i$ .                                  |
| $\alpha_{M_i}(\mathcal{P})$             | TPU weight miss probability of $M_i$ under partitioning $\mathcal{P}$ . |
| $T^{Load}$                              | Latency to load weights into TPU memory.                                |
| $s^{TPU}, s^{CPU}$                      | Service time for TPU prefix and CPU suffix.                             |
| $\mu$                                   | Service rate (1/s)  |
| $\rho$                                  | Resource utilization ( $\lambda/\mu$ )                                  |
| $E[W]$                                  | Expected queueing delay (TPU or CPU).                                   |

Let  $s_{M_{i,[a:b]}}^{TPU}$  and  $s_{M_{i,[a:b]}}^{CPU}$  denote the service time for a model block  $M_{i,[a:b]}$  on the TPU and host CPU, respectively. For TPU execution, this service time includes the deterministic compute cost and the intra-model swapping overhead incurred when the prefix footprint exceeds TPU SRAM capacity  $C$ . To reduce the search space for real-time adaptation, we restrict each model to select a single partition point. We define  $\mathcal{P} = \{p_1, \dots, p_n\}$  as the global model partitioning vector.

To model inter-model swapping overheads, we introduce a weight miss probability  $\alpha_{M_i}(\mathcal{P}) \in [0, 1]$ , representing the likelihood that a request for  $M_i$  requires a weight reload into TPU memory under global partitioning  $\mathcal{P}$ . As detailed in the following subsection,  $\alpha_{M_i}$  is approximated by aggregate TPU memory footprint and relative arrival rates of concurrent models. We denote  $T_{M_{i,[1:p_i]}}^{Load}$  as the inter-model swapping latency, which represents the time required to fetch the weights of the TPU prefix of model  $M_i$  from host memory. This latency is a deterministic value derived from the size of block  $M_{i,[1:p_i]}$  and the measured memory-to-TPU bandwidth  $B$ . Table I summarizes the notation used in our model.

Given that the TPU processes multiple model segments in a first-come, first-served (FCFS) manner [2], we model the TPU as a single, unified  $M/G/1/FCFS$  queueing system, where request arrivals follow a Poisson process and service times have a general (arbitrary) distribution. The expected wait time for such systems is well-known and is given by the Pollaczek-Khinchine formula [14]:

$$E[W^{TPU}] = \frac{\lambda^{TPU} E[(s^{TPU})^2]}{2(1 - \rho^{TPU})} \quad (1)$$

where  $\lambda^{TPU} = \sum_{i=1}^n \mathbb{I}(p_i > 0) \lambda_{M_i}$  is the aggregate arrival rate of model prefixes on the TPU, and  $\rho^{TPU} = \lambda^{TPU} E[s^{TPU}]$  is the TPU utilization. The effective service time  $s^{TPU}(\mathcal{P})$  is a weighted average of the service times for

all model prefixes under global partitioning  $\mathcal{P}$ , including the potential inter-model swapping latency:

$$E[s^{TPU}(\mathcal{P})] = \sum_{M_i \in \mathcal{M}} \frac{\lambda_{M_i}}{\lambda^{TPU}} \left( \alpha_{M_i}(\mathcal{P}) T_{M_{i,[1:p_i]}}^{Load} + s_{M_{i,[1:p_i]}}^{TPU} \right) \quad (2)$$

Following TPU processing, the suffix block  $M_{i,[p_i+1:P_i]}$  is offloaded to the host CPU. To ensure performance isolation, SwapLess partitions CPU resources by assigning each model  $M_i$  a dedicated set of  $k_i$  cores and an independent execution queue. Given that model service times are largely deterministic [2], we model the CPU suffix execution as an  $M/D/k_i$  queueing system, where  $D$  represents the deterministic service time distribution. The expected queueing delay  $E[W_{M_i}^{CPU}]$  for model  $M_i$  can be approximated as [15]:

$$E[W_{M_i}^{CPU}] = \frac{1}{2} \left( \frac{1}{k_i \mu_{M_i}^{CPU} - \lambda_{M_i}} - \frac{1}{k_i \mu_{M_i}^{CPU}} \right) \quad (3)$$

where the service rate  $\mu_{M_i}^{CPU}$  is the inverse of the service time for the CPU suffix  $s_{M_{i,[p_i+1:P_i]}}^{CPU}$ .

To model the data transfer costs between the TPU and host, let  $d_{M_i}^{in}$  denote the input size of model  $M_i$ ,  $d_{p_i}^{out}$  denote the intermediate output size at partition point  $p_i$ . Further, let  $\mathcal{K} = \{k_1, \dots, k_n\}$  denotes the core allocation for all models. The expected end-to-end latency for model  $M_i$  under a global configuration  $(\mathcal{P}, \mathcal{K})$

$$T_{M_i}^{e2e}(\mathcal{P}, \mathcal{K}) = \mathbb{I}(p_i > 0) \left[ \frac{d_{M_i}^{in}}{B} + E[W^{TPU}] + \alpha_{M_i}(\mathcal{P}) T_{M_{i,[1:p_i]}}^{Load} + s_{M_{i,[1:p_i]}}^{TPU} + \frac{d_{p_i}^{out}}{B} \right] + \mathbb{I}(p_i < P_i) \left[ E[W_{M_i}^{CPU}] + s_{M_{i,[p_i+1:P_i]}}^{CPU} \right] \quad (4)$$

**Joint Partitioning and Core Allocation.** Building on the latency model derived above, we formulate the optimization task as a Non-Linear Integer Program (NLIP). The objective is to minimize the aggregate weighted end-to-end latency across all models  $\mathcal{M}$ . Given a total physical capacity of  $K_{max}$  CPU cores, the optimization problem is defined as:

$$\min_{\mathcal{P}, \mathcal{K}} \sum_{M_i \in \mathcal{M}} \lambda_{M_i} T_{M_i}^{e2e}(\mathcal{P}, \mathcal{K}) \quad (5)$$

s.t.

$$p_i \in \{0, 1, \dots, P_i\}, \quad \forall M_i \in \mathcal{M} \quad (6)$$

$$k_i \in \{0, 1, \dots, K_{max}\}, \quad \forall M_i \in \mathcal{M} \quad (7)$$

$$\mathbb{I}(p_i < P_i) \leq k_i \leq K_{max} \cdot \mathbb{I}(p_i < P_i), \quad \forall M_i \in \mathcal{M} \quad (8)$$

$$\sum_{i=1}^n k_i \leq K_{max} \quad (9)$$

where (6) and (7) set the discrete boundaries for partition points and core allocations. (8) ensures that at least one core is allocated to any model with a CPU suffix and prevents CPU allocation for full-TPU execution. (9) enforces the global physical core capacity  $K_{max}$ .

**Modeling the Weight Miss Probability.** To model the weight miss probability  $\alpha_{M_i}$  of model  $M_i$ , let  $W(\mathcal{P}) = \sum_{M_i \in \mathcal{M}} \text{size}(M_i, [1:p_i])$  be the aggregate TPU memory footprint under partitioning  $\mathcal{P}$ . Since memory eviction policies are typically proprietary and undocumented, we adopt a conservative approximation of  $\alpha_{M_i}$  under two operational regimes:

$$\alpha_{M_i} = \begin{cases} 0 & W(\mathcal{P}) \leq C \text{ or } |\mathcal{P}| = 1 \\ 1 - \frac{\lambda_{M_i}}{\lambda_{TPU}} & \text{otherwise} \end{cases} \quad (10)$$

In the first regime, we set  $\alpha_{M_i}$  to zero because the cumulative memory requirement  $W(\mathcal{P})$  fits within TPU capacity  $C$ . Following a cold start, the system reaches a steady state where weights persist on-chip. Furthermore, for single-tenant execution ( $|\mathcal{P}| = 1$ ), we empirically observe that the reload probability remains zero regardless of model size, as the TPU runtime maintains weight persistence across inferences by fetching only required tiles rather than performing a full context reload. The second regime applies when the aggregate footprint  $W(\mathcal{P})$  exceeds  $C$ . Here, the TPU behaves as a shared-occupancy cache where  $M_i$  is subject to eviction by competing requests of other models. We approximate the residency of  $M_i$  using the stationary probability that at least one intervening request for a different model has displaced its weights since its last execution. The term  $1 - \frac{\lambda_{M_i}}{\lambda_{TPU}}$  serves as a conservative upper bound, assuming that any intervening request for a different model results in an eviction of  $M_i$ .

### C. Greedy Hill-Climbing Resource Allocation

Building on the response-time estimates from the analytic model, the optimization problem is inherently nonlinear due to the nonconvexity of the queueing models. While such problems can be addressed using standard non-linear solvers, the computational overhead is prohibitive for the resource-constrained devices targeted in this work. Therefore, we propose a lightweight heuristic that jointly optimizes model partitioning and core allocation, detailed in Algorithm 1. The algorithm first assigns all model to the CPU, with each model’s core allocation proportional to its CPU workload, and computes the expected system latency using the analytic model (Lines 1–3). In each iteration, it considers moving up to two layers to the TPU for each model  $M_m$  (Lines 6–9), performs a proportional fair-share allocation that assigns a larger integer share of the  $K_{\max}$  cores to models with higher CPU workloads (Line 10), and estimates the resulting system latency  $L'_{m,h}$  (Line 11). The search follows a greedy policy, committing the model and step  $(m^*, h^*)$  that most reduces system latency (Lines 12–16). This process repeats until no further moves reduce the expected system latency. The runtime overhead is minimal ( $< 3$  ms on the tested platform) since the algorithm relies solely on closed-form computations in the analytic models.

## IV. IMPLEMENTATION

We implement SwapLess using TensorFlow Lite, a widely deployed runtime for edge AI applications. SwapLess follows

---

### Algorithm 1 Greedy Hill-Climbing Resource Allocation

---

**Require:** Models  $\mathcal{M}$ , Rates  $\Lambda$ , Max Cores  $K_{\max}$ , TPU Capacity  $C$

**Ensure:** Optimized Partitioning  $\mathcal{P}$ , Core Allocation  $\mathcal{K}$

```

1:  $\mathcal{P} \leftarrow [p_1, \dots, p_n] = [0, \dots, 0]$ 
2:  $\mathcal{K} \leftarrow \text{PROPALLOC}(\mathcal{M}, \Lambda, \mathcal{P}, K_{\max})$ 
3:  $L_{curr} \leftarrow \sum_{M_i \in \mathcal{M}} \lambda_{M_i} T_{M_i}^{e2e}(\mathcal{P}, \mathcal{K})$ 
4: while True do
5:    $L'_{m,h} \leftarrow \infty, \mathcal{K}'_{m,h} \leftarrow \text{null}, \forall M_m \in \mathcal{M}, \forall h \in \{1, 2\}$ 
6:   for each  $M_m \in \mathcal{M}$  do
7:     for  $h \leftarrow 1$  to 2 do
8:       if  $p_m + h \leq P_m$  then
9:          $\mathcal{P}' \leftarrow [p_1, \dots, p_m + h, \dots, p_n]$ 
10:         $\mathcal{K}'_{m,h} \leftarrow \text{PROPALLOC}(\mathcal{M}, \Lambda, \mathcal{P}', K_{\max})$ 
11:         $L'_{m,h} \leftarrow \sum_{M_i \in \mathcal{M}} \lambda_{M_i} \cdot T_{M_i}^{e2e}(\mathcal{P}', \mathcal{K}'_{m,h})$ 
12:    $L_{min} \leftarrow \min_{m,h}(L'_{m,h})$ 
13:   if  $L_{min} < L_{curr}$  then
14:      $(m^*, h^*) \leftarrow \arg \min_{m,h} L'_{m,h}$ 
15:      $p_{m^*} \leftarrow p_{m^*} + h^*$ ;  $L_{curr} \leftarrow L_{min}$ ;  $\mathcal{K} \leftarrow \mathcal{K}'_{m^*,h^*}$ 
16:   else
17:     break
18: return  $\mathcal{P}, \mathcal{K}$ 

```

---

a decoupled design with an offline profiling phase and an online orchestration engine. In the offline phase, SwapLess takes pre-trained TensorFlow frozen graphs as input and performs a topological traversal to identify candidate partition points that separate the graph along a single edge. SwapLess then partitions the graphs at each candidate point into a TPU prefix and a CPU suffix using the GraphSurgeon library [16]. Segments are then converted to 8-bit quantized format via the TensorFlow Lite Converter [17], and TPU prefixes are compiled with the TPU Compiler [18] to generate optimized binaries. Finally, SwapLess performs a one-time profiling of all segments to obtain service times for its analytic model. Note that other service-time prediction methods such as [19] can be seamlessly integrated into SwapLess.

During the online phase, SwapLess continuously monitors request rates using a sliding window and periodically runs its resource allocation algorithm to update partition points and CPU core assignments. Requests are executed via a global TPU worker that maintains an FCFS queue for prefix execution. Intermediate outputs are forwarded to model-specific CPU threadpools for suffix execution, with pool sizes determined by the allocation scheme.

## V. EXPERIMENTAL EVALUATION

In this section, we present our experimental setup, validate our analytic queueing model, and compare SwapLess against three baselines across workloads and utilization levels.

### A. Experimental Setup

1) *Hardware:* Our evaluation is conducted on a Google Coral USB Accelerator featuring an Edge TPU coprocessor capable of performing 4 trillion operations per second (TOPS). The Edge TPU connects to a Raspberry Pi 5 via USB 3.0, which features a quad-core ARM Cortex-A76 CPU at 2.4 GHz and 8 GB of LPDDR4X SDRAM.

TABLE II  
CHARACTERISTICS OF EVALUATED AI MODELS.

| Model        | Size (MB) | FLOPs (G) | # Partition Points |
|--------------|-----------|-----------|--------------------|
| SqueezeNet   | 1.4       | 0.81      | 2                  |
| MobileNetV2  | 4.1       | 0.30      | 5                  |
| EfficientNet | 6.7       | 0.39      | 6                  |
| MnasNet      | 7.1       | 0.31      | 7                  |
| GPUNet       | 12.2      | 0.62      | 5                  |
| DenseNet201  | 19.7      | 4.32      | 7                  |
| ResNet50V2   | 25.3      | 4.49      | 8                  |
| Xception     | 26.1      | 8.38      | 11                 |
| InceptionV4  | 43.2      | 12.27     | 11                 |

2) *Models*: We evaluate SwapLess using nine representative AI models spanning a diverse range of computational requirements and memory footprints. Their key characteristics are summarized in Table II.

3) *Baselines*: We use three representative baselines.

- *Edge TPU Compiler* [18]: This baseline represents the current industry deployment standard. It employs a static co-compilation strategy where multiple models are compiled together to share the TPU memory. Models may incur inter-model swapping overhead during inference.
- *Threshold-based Partitioning*: This heuristic evaluates layers starting from the last one and offloads a layer to the CPU if its CPU execution time is within 10% of its TPU execution time. It focuses on per-model hardware efficiency and does not account for queuing effects or multi-tenant resource contention.
- *SwapLess* ( $\alpha = 0$ ): To isolate the impact of inter-model swapping, this variant of SwapLess excludes the weight miss probability  $\alpha_{M_i}(\mathcal{P})$ . It relies on queuing models but assumes no swapping ( $\alpha = 0$ ), thereby ignoring memory-swapping latency.

Note that *Threshold-based Partitioning* and *SwapLess* ( $\alpha = 0$ ) are baselines introduced in this work, as prior work does not provide directly comparable baselines for our single-Edge-TPU device setting.

### B. Model Validation

First, we validate the proposed analytic queuing model across single- and multi-tenant deployment scenarios.

1) *Single-tenant deployment validation*: Fig. 5 shows the accuracy of SwapLess’s analytic model using InceptionV4 under low load ( $\rho = 0.2$ ) across partition configurations. As shown, SwapLess closely matches observed execution times across configurations, achieving a mean absolute percentage error (MAPE) of 1.9%, with 92.3% of predictions within  $\pm 5\%$  and all within  $\pm 10\%$  of the observed mean latency. Fig. 5b validates the prediction accuracy across different request rates. Importantly, the optimal partition configuration depends on the request rate: PP 9 yields lower average latency below 4.5 RPS, while PP 7 becomes optimal above 4.5 RPS. This highlights the inefficiency of static configurations and motivates the need for our analytic model. Similar accuracy is observed for other models and results are omitted due to space constraints.




Fig. 5. Model validation for single AI model deployments: (a) InceptionV4 performance under different partition points (PPs). (b) InceptionV4 performance across request rates.

2) *Multi-tenant deployment validation*: Fig. 6 shows the accuracy of SwapLess’s analytic model across diverse model mixes. Fig. 6a demonstrates that the weight miss probability  $\alpha$  effectively characterizes model-loading overhead. In the first scenario, MobileNetV2 and SqueezeNet together fit within the Edge TPU memory and do not incur swapping overhead, which is correctly captured by the first regime in (10) with  $\alpha = 0$ . In the second scenario, EfficientNet and GPUNet exceed the TPU memory capacity. Here, SwapLess sets  $\alpha = 0.5$  for each model based on a 50:50 request mix, accurately modeling inference performance. In the third scenario, with a 90:10 request skew, SwapLess adjusts  $\alpha$  for each model based on the request distribution, providing precise latency predictions with a MAPE of 2.2%. Fig. 6b compares observed and predicted latencies across different model mixes, while Fig. 6c validates SwapLess’s analytic model under different request rates. For each workload mix, each model’s request rate is configured to generate an equal TPU load. The results show that SwapLess’s predicted latencies closely match the observed mean latencies for all workload mixes with a MAPE of 6.8%, confirming the accuracy of the analytic model.

**Key takeaway.** *SwapLess accurately predicts end-to-end latency across partition configurations, request rates, and workload mixes, achieving a MAPE of 1.9% for single-tenant and 6.8% for multi-tenant workloads.*

### C. Comparison with Baselines

Fig. 7 compares the performance of SwapLess with representative baselines across diverse model mixes, varying TPU utilization levels ( $\rho$ ), and both single-tenant (left) and multi-tenant (right) workloads. For multi-tenant workloads, each model contributes equally to the total request load, and the y-axis reports the mean latency across all models. As seen, when all models fit within the TPU memory (e.g., MobileNetV2 and MobileNetV2+SqueezeNet), all approaches achieve similar performance as executing the workload on TPU does not incur swapping overhead. Conversely, for workloads whose memory footprint exceeds TPU capacity, SwapLess demonstrates significant advantages. Under low utilization ( $\rho = 0.2$ ), SwapLess reduces mean latency by up to 56.2% in single-tenant settings and 68.0% in multi-tenant scenarios relative to the TPU Compiler. These substantial gains are primarily attributed to SwapLess’s explicit modeling




Fig. 6. Model validation for multi-tenant deployments: (a) Validation of the  $\alpha$  parameter across multiple workload mixes. (b) Accuracy across different concurrently running model combinations. (c) Accuracy across request rates for a model combination.

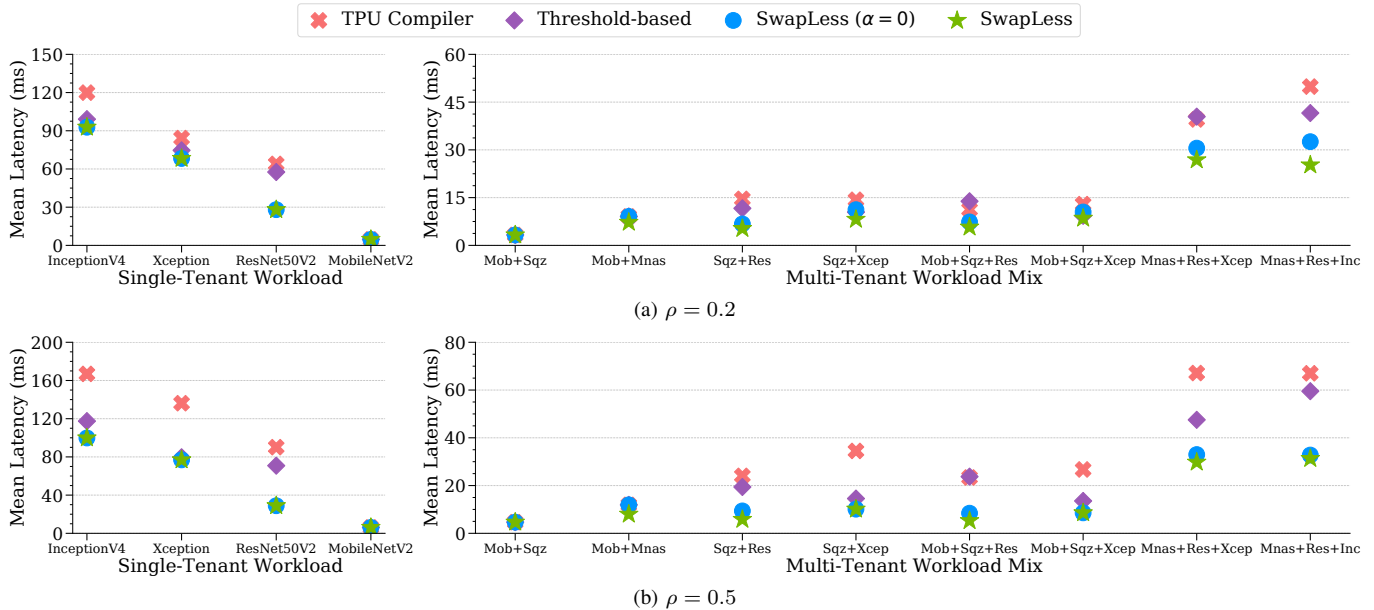


Fig. 7. Latency comparison between SwapLess and baselines under different TPU utilization  $\rho$ .

of memory swapping, which enables more effective resource allocation under memory pressure. It is worth noting that in single-tenant environments, setting  $\alpha = 0$  yields performance identical to that of the full SwapLess since no inter-model swapping occurs. As load increases to a moderate level ( $\rho = 0.5$ ), the benefits of SwapLess become even more pronounced as it accounts for both queuing delays and memory swapping overheads, achieving latency reductions of up to 63.8% and 77.4% in single- and multi-tenant settings, respectively. In contrast, threshold-based partitioning performs worse than the TPU compiler for some workloads (e.g., MobileNetV2+SqueezeNet+ResNet), as it does not account for queuing delays and swapping overheads. Overall, SwapLess consistently matches or outperforms all baselines, achieving the lowest overall latency.

**Key takeaway.** SwapLess reduces end-to-end latency by up to 63.8% and 77.4% compared to the baselines across both single- and multi-tenant deployments.

#### D. Impact of Dynamic Workloads

We now evaluate SwapLess’s ability to sustain low latency under time-varying load. Fig. 8 shows the performance of




Fig. 8. Performance under dynamic request rates for MnasNet and InceptionV4. Request rates are initially (5,1) RPS, then (5,3) from 300–600s, and (5,5) from 600–900s.

SwapLess and representative baselines as workload request rates vary over time. To enable fast reconfiguration, SwapLess preloads a small set of representative model partitions for low-overhead configuration switching, with further optimization of the switching mechanism left for future work. Initially, SwapLess partitions TPU memory to accommodate both models, with MnasNet fully running on the TPU and InceptionV4 partitioned at  $p_i = 7$ . As InceptionV4’s request

rate increases to 3 RPS, SwapLess dynamically offloads the MnasNet model to the CPU, reallocating the freed TPU resources to InceptionV4 to handle the increased load. When the load further increases to 5 RPS, SwapLess offloads one additional InceptionV4 layer to CPU to mitigate TPU queuing delays and reallocates an additional core to handle its increased CPU load. The overhead of the adaptive algorithm is negligible and only incurs less than 3 ms per invocation. Through adaptive resource allocation, SwapLess achieves up to 75.1% latency reduction compared to baselines.

**Key takeaway.** *SwapLess adapts to workload changes with minimal overhead, reducing mean latency by up to 75.1%.*

## VI. RELATED WORK

**Edge TPU Inference Optimization.** Several work has attempted to address the memory constraints on Edge TPUs [8]–[11], [20], [21]. For example, RESPECT [9] proposes a reinforcement learning–based DNN scheduler for pipelined Edge TPUs. SAPar [21] introduces a surrogate-assisted partitioning tool to distribute DNN segments across multiple Edge TPUs to reduce latency. However, these approaches rely on multiple TPU devices and do not address the challenges of single-device execution, which is the focus of our work.

**Cross-processor Inference.** Prior work has explored partitioning DNNs across processors such as CPUs and GPUs to improve inference performance [22]–[26]. For example, Band [25] partitions and schedules model subgraphs across mobile processors, while BlastNet [24] partitions DNNs into paired CPU-GPU blocks across processors. However, they assume a unified memory architecture in which accelerators can directly access host memory without incurring swapping overhead, and thus do not address the challenges introduced by memory-constrained accelerators.

**Multi-tenant Inference.** Another line of work focuses on coordinating and optimizing the execution of multiple concurrent DNNs [27]–[30]. For instance, LegoDNN [29] uses block-grained scaling to dynamically combine reusable blocks at runtime, optimizing accuracy and latency. POS [30] employs a reinforcement learning–based operator scheduler to coordinate concurrent DNN execution to reduce latency. However, these approaches generally assume sufficient memory and computing resources, and do not handle the challenges introduced by memory-constrained accelerators.

## VII. CONCLUSION

This paper presented SwapLess, a system for multi-tenant TPU–CPU collaborative inference on memory-constrained Edge TPU devices. SwapLess uses an analytic queuing model that captures collaborative inference performance characteristics for adaptive resource allocation. Evaluation results show that SwapLess reduces mean inference latency by up to 77.4% compared to the default Edge TPU compiler.

## ACKNOWLEDGMENT

This research is supported by NSF grants 2211302, 2211888, 2213636, 2105494, 23091241, 2325956, and the

Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (IoBT CRA).

## REFERENCES

- [1] X. Wang *et al.*, “Empowering Edge Intelligence: A Comprehensive Survey on On-Device AI Models,” *ACM Comput. Surv.*, vol. 57, Apr. 2025.
- [2] Q. Liang *et al.*, “Model-driven Cluster Resource Management for AI Workloads in Edge Clouds,” *ACM Trans. Auton. Adapt. Syst.*, vol. 18, Mar. 2023.
- [3] N. Ng *et al.*, “Collaborative Inference in Resource-Constrained Edge Networks: Challenges and Opportunities,” in *MILCOM’24*, 2024.
- [4] T. Abdelzaher *et al.*, “Will Distributed Computing Revolutionize Peace? The Emergence of Battlefield IoT,” in *JCDCS*, pp. 1129–1138, 2018.
- [5] Q. Dong, J. Xu, P. Pillai, and M. Satyanarayanan, “Does Accurate Real-Time AI Need Edge Offload?,” in *SEC ’25*, 2025.
- [6] Google, “Edge TPU: Run Inference at the Edge.” <https://www.coral.ai/docs/edgetpu/inference/>, 2019. Accessed: 2026-01-23.
- [7] “Raspberry Pi AI HAT+.” <https://www.raspberrypi.com/products/ai-hat/>. Accessed: 2026-01-23.
- [8] B. Zou *et al.*, “A Performance Prediction-based DNN Partitioner for Edge TPU Pipelining,” in *MILCOM 2024*, pp. 1–6, 2024.
- [9] J. Yin *et al.*, “RESPECT: Reinforcement Learning Based Edge Scheduling on Pipelined Coral Edge TPUs,” in *DAC’23*, p. 1–6, 2025.
- [10] J. Villarrubia *et al.*, “Improving inference time in multi-TPU systems with profiled model segmentation,” in *PDP’23*, pp. 84–91, 2023.
- [11] J. Yin *et al.*, “Exact Memory- and Communication-aware Scheduling of DNNs on Pipelined Edge TPUs,” in *SEC’22*, pp. 203–215, 2022.
- [12] H. Li *et al.*, “Enabling Real-time AI Inference on Mobile Devices via GPU-CPU Collaborative Execution,” in *RTCSA’22*, 2022.
- [13] B. Reidy *et al.*, “Work in Progress: Real-time Transformer Inference on Edge AI Accelerators,” in *RTAS’23*, pp. 341–344, 2023.
- [14] M. Harchol-Balter, *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press, 2013.
- [15] A. Shah *et al.*, “Adaptive Alert Management for Balancing Optimal Performance among Distributed CSOCs using Reinforcement Learning,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 31, p. 16–33, Jan. 2020.
- [16] NVIDIA, “ONNX GraphSurgeon.” <https://pypi.org/project/onnx-graphsurgeon/>, 2025. Accessed: 2026-01-20.
- [17] “tf.lite.TFLiteConverter.” [https://www.tensorflow.org/api\\_docs/python/tf/lite/TFLiteConverter](https://www.tensorflow.org/api_docs/python/tf/lite/TFLiteConverter), 2024. Accessed: 2026-01-21.
- [18] Coral AI, “Edge TPU Compiler.” <https://www.coral.ai/docs/edgetpu/compiler>. Accessed: 2026-01-19.
- [19] S. Yao *et al.*, “FastDeepIoT: Towards Understanding and Optimizing Neural Network Execution Time on Mobile and Embedded Devices,” in *SenSys’18*, 2018.
- [20] “Pipeline a Model with Multiple Edge TPUs.” <https://gweb-coral-full.ucr.appspot.com/docs/edgetpu/pipeline/>. Accessed: 2026-01-21.
- [21] B. Sun *et al.*, “Sapar: A surrogate-assisted dnn partitioner for efficient inferences on edge tpu pipelines,” *ACM Trans. Embed. Comput. Syst.*, vol. 24, Sept. 2025.
- [22] Y. Kim *et al.*, “ $\mu$ layer: Low latency on-device inference using cooperative single-layer acceleration and processor-friendly quantization,” in *EuroSys’19*, 2019.
- [23] F. Jia *et al.*, “CoDL: efficient CPU-GPU co-execution for deep learning inference on mobile devices,” in *MobiSys’22*, p. 209–221, 2022.
- [24] N. Ling *et al.*, “BlastNet: Exploiting Duo-Blocks for Cross-Processor Real-Time DNN Inference,” in *SenSys’22*, p. 91–105, 2023.
- [25] J. S. Jeong *et al.*, “Band: coordinated multi-dnn inference on heterogeneous mobile processors,” in *MobiSys’22*, p. 235–247, 2022.
- [26] T. Sen *et al.*, “Flex: Fast, Accurate DNN Inference on Low-Cost Edges Using Heterogeneous Accelerator Execution,” in *EuroSys’25*, p. 507–523, 2025.
- [27] A. H. Jiang *et al.*, “Mainstream: Dynamic Stem-Sharing for Multi-Tenant video processing,” in *USENIX ATC’18*, pp. 29–42, July 2018.
- [28] B. Fang *et al.*, “NestDNN: Resource-Aware Multi-Tenant On-Device Deep Learning for Continuous Mobile Vision,” in *MobiCom’18*, 2018.
- [29] R. Han *et al.*, “LegoDNN: block-grained scaling of deep neural networks for mobile vision,” in *MobiCom’21*, p. 406–419, 2021.
- [30] Z. Zhang *et al.*, “POS: An Operator Scheduling Framework for Multi-model Inference on Edge Intelligent Computing,” in *IPSN’23*, 2023.