RecPipe: Co-designing Models and Hardware to Jointly Optimize Recommendation Quality and Performance

Udit Gupta[1,2], Samuel Hsia[1], Jeff (Jun) Zhang[1], Mark Wilkening[1], Javin Pombra[1],
Hsien-Hsin S. Lee[2], Gu-Yeon Wei[1], Carole-Jean Wu[2], David Brooks[1]

[1]Harvard University, [2]Facebook AI Research

ugupta@g.harvard.edu

ABSTRACT

Deep learning recommendation systems must provide high quality, personalized content under strict tail-latency targets and high system loads. This paper presents RecPipe, a system to jointly optimize recommendation quality and inference performance. Central to RecPipe is decomposing recommendation models into multi-stage pipelines to maintain quality while reducing compute complexity and exposing distinct parallelism opportunities. RecPipe implements an inference scheduler to map multi-stage recommendation engines onto commodity, heterogeneous platforms (e.g., CPUs, GPUs). While the hardware-aware scheduling improves ranking efficiency, the commodity platforms suffer from many limitations requiring specialized hardware. Thus, we design RecPipeAccel (RPAccel), a custom accelerator that jointly optimizes quality, tail-latency, and system throughput. RPAccel is designed specifically to exploit the distinct design space opened via RecPipe. In particular, RPAccel processes queries in sub-batches to pipeline recommendation stages, implements dual static and dynamic embedding caches, a set of top-k filtering units, and a reconfigurable systolic array. Compared to prior-art and at iso-quality, we demonstrate that RPAccel improves latency and throughput by $3\times$ and $6\times$.

1. Introduction

Deep neural network (DNN) based recommendation systems constitute an overwhelming fraction of AI cycles in production data centers (e.g., Facebook, Google, Alibaba) [1, 18, 21, 22, 30, 51, 58, 59, 60]. To improve content personalization in a wide range of services (e.g., search, e-commerce, movie and video-streaming, social media), the size of production recommendation models has grown by over $10\times$ between 2017 and 2020 [39, 56, 57].

In response to the dramatic increase in infrastructure demands from the ever-increasing model complexity, system- and architecture-level solutions are customized for DNN-based recommendation, including inference schedulers [16], near memory processing hardware [33, 36, 49], and specialized accelerators [9, 26, 29]. These prior solutions assume fixed models, leaving significant room for efficiency optimization. Evidenced by recent work optimizing DNNs for computer vision and natural language processing [17, 19, 24,
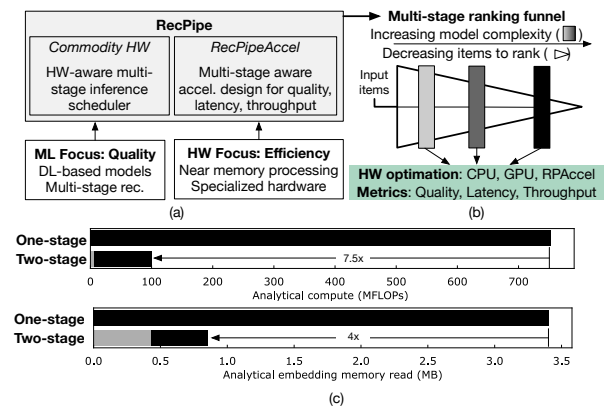

Figure 1: **(a) Compared to prior work from machine learning and hardware researchers, this work jointly optimizes quality and performance. (b) RecPipe co-designs models and hardware across multi-stage recommendation pipelines. (c) Transforming monolithic models into multiple stages reduces overall compute demand and embedding memory accesses by $7.5\times$ and $4.0\times$, respectively.**

43, 45, 47, 53], co-designing models with hardware is an effective approach. However, the model accuracy requirement for recommendation tasks is stringent [59, 60], making the model-hardware co-design space challenging to navigate.

While accuracy represents a model's ability to predict whether users will like *individual* items, production services are designed to serve users a personalized *collection* of relevant items [5, 28]. As such, while accuracy is intrinsic to models, *quality* is optimized by improving model accuracy *and* increasing the number of items ranked at the same time. The more holistic, application-level quality objective allows system architects to judiciously trade off accuracy for performance, opening new design spaces for system optimization.

Ranking all items with complex models is wasteful—only a small portion of items are relevant to individual users. Traditionally, recommendation engines achieve high quality by ranking a large number of input candidate items using complex DNNs. The combination of large input working set sizes and complex models incurs high performance overheads. Alternatively, one can decompose a monolithic ranking model

into multiple stages to maintain overall quality at higher performance [27, 32, 58]. By splitting the monolithic model into two, a frontend model coarsely *filters* relevant items while a more accurate backend model finely *ranks* items to serve. Further segmenting the pipeline into additional stages creates a ranking funnel (Figure 1 (b)) where complex models only rank items requiring accurate differentiation. For the Criteo dataset and Facebook's Deep Learning Recommendation Model (DLRM) [31, 41], Figure 1(c) shows that, at iso-quality, compared to single-stage, multi-stage recommendation reduces memory and compute demands by 4.0× and 7.5×, respectively. This system-level view optimizing quality and efficiency motivates a new generation of hardware solutions for multi-stage recommendation.

Driven by this motivation, we propose *RecPipe*, a system to co-design recommendation models and hardware to improve both quality and performance (Figure 1(a)). Frontend stages pair light-weight models (e.g., low compute and memory demands) with large input sizes, exposing *data-parallelism*. Backend stages pair heavy-weight models (e.g., billions of FLOPs, many GBs of storage) with small input sizes, exposing *model-parallelism* instead. RecPipe's system solutions exploit these distinct parallelism opportunities to jointly optimize quality, throughput, and tail-latency.

To understand the limits of commodity platforms, RecPipe implements an inference scheduler that maps each recommendation stage across heterogeneous hardware (e.g., CPU, GPU) to maximize performance. We find the optimal mapping depends on the application level targets and underlying hardware. Despite the tight co-design between models and hardware, we find commodity CPU-GPU systems do not fully exploit the benefits of multi-stage recommendation as they suffer from low utilization and high PCIe communication overheads between stages.

To address these limitations, we design *RecPipeAccel* (RPAccel), a specialized accelerator for multi-stage recommendation. Starting with a TPU-like, systolic array-based, recommendation accelerator [26], RPAccel's hardware optimizations improve efficiency at low area and power overheads. First, RPAccel implements a reconfigurable systolic array that allows the hardware to concurrently process models across recommendation stages. RecPipe's inference scheduler provisions the fraction of systolic array resources to devote to frontend and backend stages based on application load, balancing latency and throughput. Next, RPAccel eliminates high PCIe communication overheads to the host processor by implementing multiple on-chip filtering units to identify the top-*k* user-item interactions between stages. Finally, to overlap frontend and backend query processing, RPAccel breaks queries into sub-batches to pipeline stages and prefetch embedding vectors in separate caches.

The main contributions of this work include:

1. We propose a new system, RecPipe, that enables design space exploration and optimization for multi-stage recommendation inference. The framework integrates data sets (e.g., MovieLens [20], Criteo [31]), models (e.g., neural matrix factorization [23], DLRM [41]), and hardware (e.g., CPU, GPU, simulated accelerators) to study tradeoffs among quality, tail-latency, and throughput.
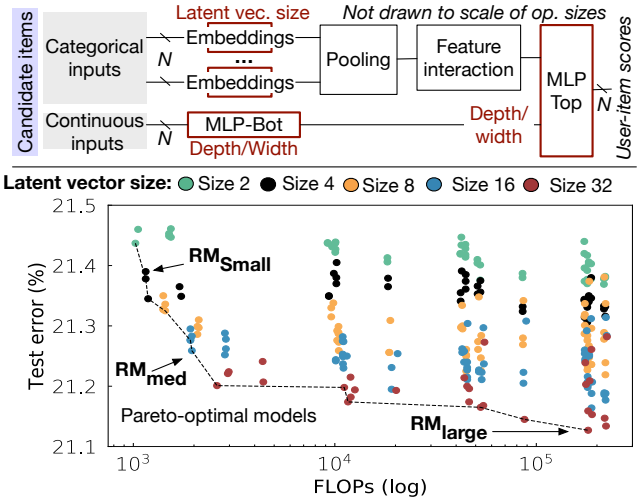

Figure 2: (Top) General recommendation model architecture configured by embedding dimension and MLP size (outlined in red). (Bottom) Hyperparameter sweep shows tradeoff between model complexity and error.

Model	RM_{small}	RM_{med}	RM_{large}
Embedding Dim.	4	16	32
MLP-Bottom	13-64-4	13-64-16	13-512-256-128-64-32
MLP-Top	64-1	64-1	96-1
Model Size	1GB	4GB	8GB
FLOPs	1.1K	2.0K	180K
Model Error	21.36%	21.26%	21.13%

Table 1: Pareto-optimal recommendation models.

2. We show designing and efficiently scheduling multi-stage pipelines for available commodity hardware platform reduces tail-latency by 4× and 3× on CPUs and heterogeneous CPU-GPU hardware, respectively.

3. We design RPAccel, a novel accelerator that exploits the distinct properties of multi-stage recommendation to jointly optimize quality, latency, and throughput. Compared to a state-of-the-art baseline accelerator [26], RPAccel's software and hardware optimizations reduce tail-latency by 3× and increases throughput by 6×, at iso-quality as well as negligible area and power overheads.

2. Motivation: Widening Design Space by Optimizing for Quality over Accuracy Alone

Prior work on specialized systems for deep learning co-optimizes for model accuracy and run-time efficiency (performance, power, and energy) [17, 19, 24, 43, 45, 47]. For neural recommendation however, hardware designers must go one step further, beyond accuracy, and optimize for quality. In this section, we first describe recommendation model architectures and conduct a model hyper-parameter sweep. Then, we introduce the quality metric used in this work, showing the fundamental difference between accuracy and quality.

2.1 Training hyperparameter sweep

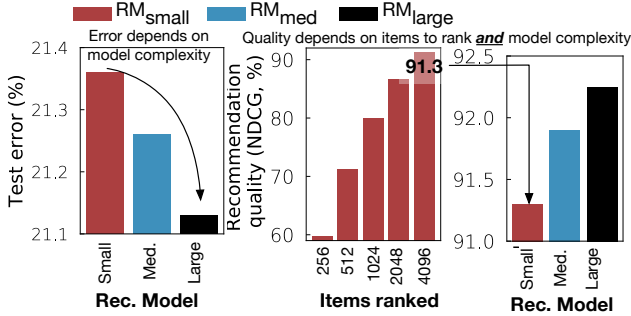Figure 2(top) lays out the general architecture for DNN recommendation models [18, 41]. Continuous input features

Figure 3: While accuracy depends only on model size (left), recommendation quality depends on number of items ranked (center) *and* model size (right).

are processed with DNN layers, e.g. Multi Layer Perceptrons (MLP), while sparse input features are processed using embedding tables. Embedding tables are organized as a collection of embedding vectors with tens to hundreds of latent features. Latent features map sparse inputs to low-dimensional, dense ones. By configuring the main network components (i.e., MLP depth/width, embedding latent vector dimension), highlighted in red, we realize models with varying storage capacity, compute demands, and accuracy.

Figure 2(bottom) shows a hyperparameter sweep by tuning the main network parameters for Facebook's DLRM trained on the Criteo dataset [31, 41]. Increasing the computational complexity of models reduces the test error. Models with 1.1K FLOPs and 180K FLOPs observe an error of 21.36% and 21.13%, respectively. Note, a 0.23% decrease in error is large given the high sparsity of user-item interactions in recommendation use cases [59, 60]. Recent industry publications show reductions of even 0.1% error greatly improve user experience in real world applications [59, 60]. Table 1 shows the tradeoff between model error and complexity across three Pareto-optimal networks (i.e., RM_{small}, RM_{med}, RM_{large}).

2.2 Quality versus accuracy

A model's accuracy represents its ability to correctly predict a user will positively interact with a *single* item. However, in recommendation, models rank thousands of items opening the door for measuring overall quality. Quality measures the relevance of the *entire collection* of items presented to users based on their personal preferences. Following recent work from machine learning and recommendation systems researchers, we use normalized discounted cumulative gain (NDCG) to quantify the quality of the ordered list of output items. NDCG [5, 28] is the ratio between the measured and the ideal ordering, each of which is computed using discounted cumulative gain (DCG): for a list of N items, $DCG = \sum_i^N \frac{Rel_i}{log_2(i+1)}$. Rel_i represents item i's score in the measured or ideal list and $log_2(i+1)$ discounts its relevance—dividing the score by the item's position in the list.

Widening design space. Compared to accuracy, optimizing for quality opens new system design opportunities. For the Criteo dataset, Figure 3 illustrates the impact of varying the number of items ranked (x-axis) and model architecture (i.e., RM_{small}, RM_{med}, RM_{large}) on quality. Empirically, the improvement in quality from increasing number of items ranked overshadows that from larger, more accurate models.

Note, the highest quality of 92.25 can be achieved by ranking all 4096 items with RM_{large}. However, as shown in Figure 1, decomposing monolithic models into multiple stages, where small models filter relevant items and large models perform fine-grained ranking, improves computational efficiency at iso-quality. At the frontend, candidate items are coarsely ranked with models that incur memory and compute demands. This reduces the list of candidate items (i.e., working set size) incrementally over the stages. Subsequent stages use larger models for finer-grained ranking. Going beyond accuracy, quality depends on the number of stages, network architectures, and the number of items ranked at run-time: widening the design space to co-optimize performance and quality.

3. RecPipe Design: A System to Optimize Multi-Stage Recommendation Inference

We propose RecPipe, a novel system to explore the model- and hardware-level design space to collectively optimize recommendation quality, tail-latency, and system throughput. Figure 4(left) shows RecPipe's multi-step design process. The input to RecPipe is a Pareto-frontier of recommendation models balancing model accuracy and complexity. To co-optimize quality and hardware efficiency on commodity platforms, RecPipe balances multi-stage parameters and statically schedules each stage across available hardware resources (i.e., CPUs and GPUs). Going further, RecPipe exposes distinct parallelism opportunities that are exploited by designing specialized hardware. Figure 4(right) illustrates the multi-stage recommendation pipeline and the design space optimized by RecPipe. The model-level and hardware-level design parameters are highlighted in red. We detail how RecPipe co-designs these parameters to maximize quality and performance below.

3.1 Hardware-aware multi-stage scheduling

RecPipe implements a post-training, inference scheduler customized for multi-stage recommendation. In step 1, RecPipe balances multi-stage modeling parameters. In step 2, RecPipe exploits the parallelism opportunities exposed from step 1, and maps stages across heterogeneous hardware.

Algorithmic scaling (Step 1). RecPipe exhaustively explores the design space of pairing Pareto-optimal recommendation models and number of items to rank at each stage in the multi-stage pipeline. In the frontend, lightweight models are paired with large working set sizes exhibiting high data-level parallelism; in the backend heavyweight models are paired with smaller working set sizes exhibiting high model-level parallelism. By collectively balancing model complexity and input working set size, RecPipe maximizes overall quality under strict latency targets and system loads.

Heterogeneous hardware mapping (Step 2). Given the distinct parallelism opportunities from the aforementioned algorithmic scaling step, RecPipe exhaustively explores the mapping of multi-stage models on available hardware at the stage granularity. We begin by considering commodity hardware platforms i.e., CPUs and GPUs. GPUs implement a highly data-parallel architecture that parallelize individual queries, especially in the frontend with large working set sizes. On the other hand, many-core CPUs can simultaneously process multiple queries providing high-throughput. RecPipe exploits these architectural differences to schedule
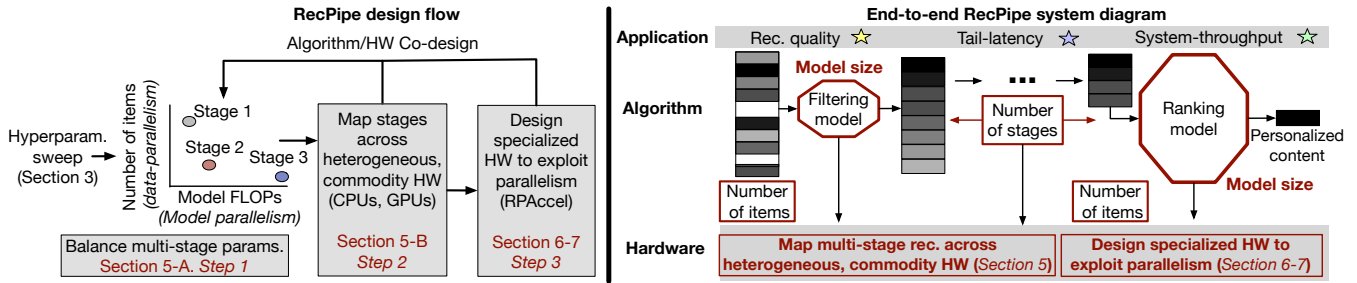
Figure 4: The structure of a multi-stage recommendation pipeline. Highlighted in red, RecPipe explores a variety of recommendation model and hardware infrastructure parameters to balance quality, latency, and throughput.
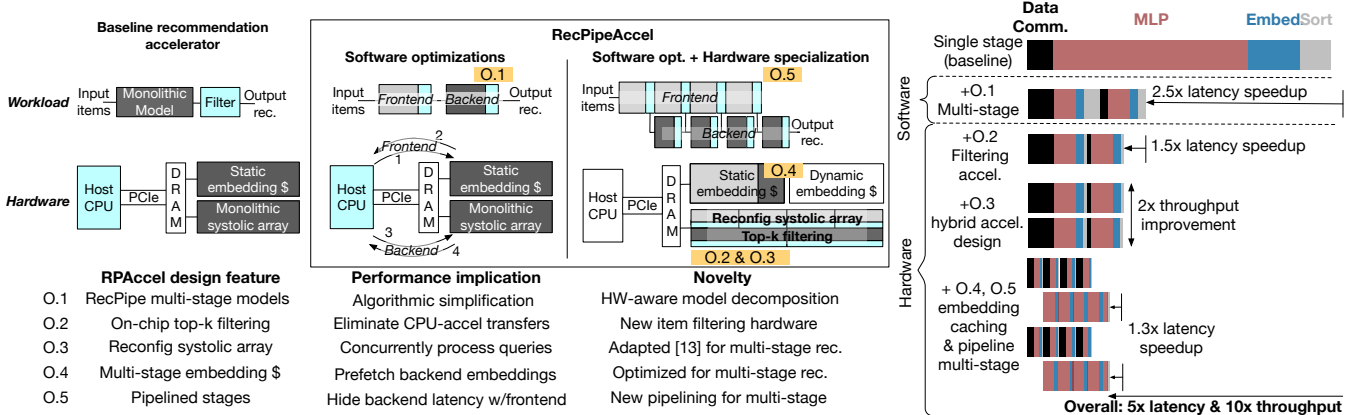

Figure 5: Comparison of baseline recommendation hardware accelerator and RPAccel. We describe RPAccel's five main innovations (i.e., O.1 to O.5) on the left and their performance benefits in the ablation study on the right.

each recommendation stage onto the underlying hardware. In fact, we find the optimal mapping of multi-stage recommendation varies across application-level targets (e.g., tail-latency, system load). Thus, RecPipe schedules multi-stage pipelines onto available hardware, based on algorithmic model parameters, architectural characteristics, and application-level requirements, to maximize quality and performance.

While achieving the maximal quality target and at iso-throughput, the scheduling optimizations reduce tail-latency by $4\times$ on CPUs and $3\times$ on heterogeneous hardware i.e., CPUs and GPUs (see Section 5 for details). However, despite these performance improvements there remains significant room for further optimization. In particular, the commodity CPU-GPU platforms suffer from two main drawbacks. First, GPUs exhibit low utilization when exploiting data-level parallelism in the frontend and model-level parallelism in the backend, primarily due to the high overhead of embedding lookups and memory transformation operations on GPUs [16]. Second, between stages, high PCIe communication overheads across the CPU and GPU limit achievable throughput. To address these limitations, and given the importance of data center-scale recommendation, RecPipe enables designing specialized hardware for multi-stage recommendation.

3.2 Custom hardware to accelerate multi-stage recommendation

Figure 5 illustrates the high-level architecture of the proposed recommendation accelerator, RPAccel. On the left, we start with a state-of-the-art accelerator baseline that minimizes inference latency for a single-stage recommendation

model using a TPU-like monolithic systolic array and static cache for *hot-embeddings* [26]. The aforementioned software optimizations reduce workload complexity by decomposing the single-stage model into a multi-stage pipeline. Given the simplified workload, RPAccel is designed to concurrently process multiple models and queries, end-to-end. Figure 5(right) provides an ablation study for the proposed software and hardware optimizations, demonstrating significant latency and throughput improvement potential (i.e., O.1 to O.5).

By exploiting unique properties of multi-stage recommendation, RPAccel is designed to balance both inference latency and throughput based on application-level requirements.

- (O.1) RecPipe decomposes a single-stage model into multiple stages ($2.5\times$ latency reduction).
- (O.2) RPAccel comprises a top-k filtering unit to identify the k highest quality items based on predicted click-through-rate (CTR) to be ranked by subsequent stages; this eliminates host-accelerator communication between recommendation stages ($1.5\times$ latency reduction).
- (O.3) RPAccel implements a reconfigurable systolic array to concurrently process multiple stages and queries ($2\times$ hardware utilization and throughput). RecPipe's software scheduler (see above) splits the monolithic systolic array into multiple sub-arrays based on application-level targets (quality, latency, throughput) and multi-stage models.
- (O.4) RPAccel balances on-chip memory resources to statically cache *hot-embeddings* and dynamically prefetch embeddings for backend models (40% reduction in average memory access time). The static cache is provisioned for both frontend and backend stages; the dynamic cache
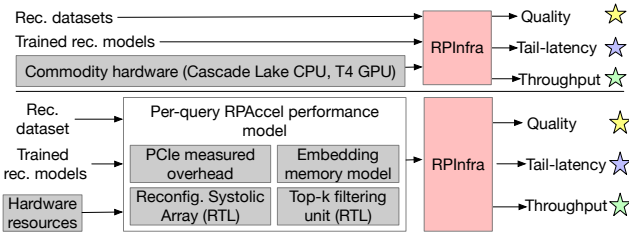
Figure 6: Evaluation methodology of RecPipe on commodity (top) and specialized (bottom) hardware.

Machines	Cascade Lake CPU	NVIDIA T4 GPU
Frequency	2.8 GHz	585 MHz
Cores	64	2560
SIMD	AVX-512	FP32x64
Cache Sizes	1-16-22 MB	96-512 KB
DRAM Capacity	384 GB	15 GB
DDR Bandwidth	75 GB/s	300 GB/s
TDP	300 Watt	70 Watt

Table 2: Commodity hardware in experimental setup.

prefetches embeddings for the backend as the frontend finishes sub-batches of the input query.

- (O.5) RPAccel breaks queries into sub-batches to pipeline – and thus – overlap computation from frontend and backend stages ($1.3\times$ latency reduction).

While achieving the highest quality target, compared to the baseline recommendation accelerator, RPAccel's optimizations collectively decrease tail-latency by up to $5\times$ and increase throughput by up to $10\times$ (see Section 6-7 for details).

4. Experimental Methodology

Figure 6 illustrates the evaluation methodology we use to study the system design implications of multi-stage recommendation. RecPipe encompasses a vast design space across multi-stage modeling parameters, hardware solutions, and application-level targets. To foster deeper understanding, we analyze cross-sections of the design space based on the application-level targets: iso-quality, iso-throughput, and iso-latency. This section details the methodology on both real, commodity hardware and simulated, specialized hardware.

Datasets and models. We evaluate RecPipe with three open-source datasets: Criteo Kaggle [31], MovieLens 1M [20], and MovieLens 20M [20]. We train neural matrix factorization models for both MovieLens datasets [22]. To provide intuition across the large design space studied in this work, we conduct a deep dive using Criteo and Facebook's DLRM [41]. On top of this deep dive, Section 8 summarizes results across all datasets. All models are implemented in PyTorch.

Application-level targets. This work optimizes recommendation based on three application-level targets:

- Quality: We use NDCG [5, 28] to quantify recommendation quality of the top sixty-four items served. For commensurate analysis, final results are presented based on the highest quality achieved for each model and dataset: NDCG of 92.25 for Criteo (see Section 2).
- Tail-latency: To maintain user-experience, recommendations must meet SLAs and be served under strict tail-latency targets [16], measured as 99^{th} percentile (p99).

Parameter	RPAccel configuration
Frequency	250 MHz
Systolic Array SRAM capacity	8MB
Systolic Array MAC units	128×128 MACs
Embedding cache capacity	16MB
DRAM capacity	16 GB
DRAM bandwidth	64 GB/s
DRAM latency	100 cycles

Table 3: Fixed resources in RPAccel.

- Throughput: Data-center recommendation systems must maximize throughput, measured as the queries processed per second (QPS). Queries follow a Poisson arrival rate.

Commodity hardware systems. To study the proposed designs in the context of data center scale recommendation, RecPipe runs datasets and models directly on real CPUs (server class Intel Cascade Lake) and GPUs (NVIDIA T4). Refer to Table 2 for detailed hardware specifications. Experiments on CPUs use multiple processes to exploit parallelism across cores—each core has a single PyTorch/MKL thread. GPUs use CUDA/cuDNN 10.1.

Accelerator modeling. RecPipe uses a two-step evaluation methodology to simulate specialized hardware.

First, we evaluate the latency of each query across each stage of the multi-stage pipeline. The latency per stage is computed as cumulative time from data transfers over PCIe, embedding lookups, MLP operations, and the top-k filtering units. Host-to-accelerator PCIe overheads are based on real measurements from the CPU-GPU system (see Table 2). For embedding lookups, we compute hit rates based on the cache locality of open-source datasets. Given the cache hit rates, we compute the memory latency of embedding operations using simple latency and bandwidth models for SRAM and DRAM. For MLP layers, We design and implement the systolic array and the top-k filtering unit in RTL to gather cycle-accurate performance measurements, including overheads from loading weights and activations from DRAM. Combining latency for all stages forms per-query performance model.

Second, the per-query latencies are fed into RecPipe which simulates the at-scale performance characteristics of RPAccel, measuring tail-latency, system-throughput, and quality, of processing tens of thousands of queries.

For area and power evaluations, we separately synthesize the reconfigurable systolic array, top-k filtering unit, and memories in a $12nm$ FinFET technology. As shown in Table 3, RPAccel implements comparable compute and memory resources to a data center TPU accelerator (40 Watt TDP) [30].

5. Evaluation of RecPipe Inference Scheduler on Commodity Hardware

In this section we use RecPipe to efficiently schedule multi-stage recommendation onto heterogeneous hardware available in data centers. First, RecPipe balances the multi-stage modeling parameters—number of stages, models per stage, items to rank per stage—to co-optimize tail-latency, throughput, and quality. Next, RecPipe co-designs the multi-stage parameters for heterogeneous systems comprising CPUs and GPUs. We show the optimal configuration of multi-stage parameters depends on the underlying hardware. Furthermore, we show that while GPUs enable higher throughput and
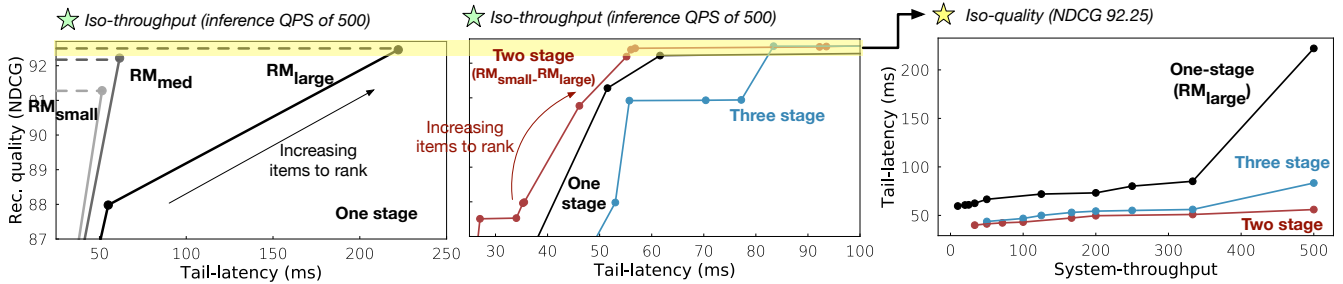
Figure 7: **(Left) In single-stage recommendation, larger models achieve the higher quality at the expense of tail-latency. (Middle) Tuning multi-stage parameters improves quality under strict performance constraints. (Right) While achieving the highest-quality target, decomposing single-stage recommendation to multiple stages reduces tail-latency.**

quality at low-latency targets, CPU-only execution achieves higher throughput under more relaxed latency targets.

5.1 Mapping multi-stage pipelines to CPUs

Figure 7(left) illustrates the tradeoff between tail-latency and quality for single-stage recommendation on CPUs. Following intuition, larger more complex models (e.g., RM_{large}) achieve higher quality at the expense of higher tail-latency.

Takeaway 1: *Carefully balancing multi-stage parameters unlocks higher recommendation quality and throughput at strict tail-latency targets.*

At a fixed system load (i.e., QPS of 500), Figure 7(center) shows tradeoff between tail-latency and quality for one-, two-, and three-stage designs. Exhaustively sweeping all possible combinations of models per stage and number of items to rank per stage, we show the Pareto-frontier results.

Compared to single-stage designs, Figure 7(center) shows multi-stage designs achieve higher quality under strict performance constraints. The single-stage design ranks all 4096 items with RM_{large}. The optimal two-stage design first processes 4096 items with RM_{small} followed by the top 256 items with RM_{large}, reducing tail-latency by $4\times$ given the lower compute and memory demands.

The importance of optimizing for quality, not accuracy, can be seen by diving deeper into the two-stage design. To achieve high quality, the backend implements the most accurate network (i.e., RM_{large}); the frontend implements either RM_{med} or RM_{small}. While RM_{med} has higher accuracy, the benefits are overshadowed by the additional compute and memory requirements (see Table 1). In fact, with RM_{large} in the backend, both frontend options achieve the same quality (NDCG 92.25). But, the combination of RM_{med}-RM_{large} has a $1.6\times$ longer tail-latency compared to RM_{small}-RM_{large}. Designers must jointly optimize for quality and performance.

In addition to quality, balancing multi-stage parameters improves throughput at strict tail-latency targets. Figure 7(right) shows the tradeoff between tail-latency and throughput, at the highest quality target (NDCG of 92.25). Compared with the one-stage system, the two-stage pipeline reduces tail-latency by $4.4\times$ (QPS of 500). However, decomposing the pipeline into three stages decreases performance given additional queuing delays between stages, which overshadow the 30% reduction in compute between two- and three-stage designs. Note, the tradeoffs will vary across datasets—varying model complexities and items to rank per stage will impact the optimal configuration (see Section 8 for examples).

5.2 Mapping multi-stage pipelines to heterogeneous systems

Figure 8(top) illustrates the tradeoff between throughput and tail-latency while achieving the high quality target (NDCG of 92.25). Using RecPipe, we exhaustively evaluate all mappings between multi-stage recommendation and heterogeneous hardware and show the best configurations: one-stage GPU-only, two-stage GPU-CPU, and the two-stage CPU-only configurations in Figure 8(top). For the two-stage GPU-CPU design, RecPipe maps either the frontend or the backend to the GPU, running the other on the CPU. In particular, we show results for frontend running on the GPU and backend on the CPU as our empirical evaluations show it provides higher performance. We also evaluate mapping two stages to the GPU with multi-tenant execution. Our evaluations show this configuration is unable to extract the fine-grain parallelism from multi-stage's data dependency, incur longer latency than the one-stage GPU-only configuration.

Takeaway 2: *Given architectural differences, the optimal multi-stage parameters vary on CPUs versus GPUs.*

Recall from our previous analysis, for CPU-only execution the two-stage design achieves the highest performance; on the heterogeneous system, however, the single-stage GPU-only configuration (solid black) achieves higher performance than multi-stage using both CPU and GPU (solid red). The reason is twofold. First, we observe comparable latency for RM_{small} versus RM_{large} on the GPU, overshadowing the benefits of decomposing models into finer-grained pipelines. Second, the multi-stage GPU-CPU design requires transferring more intermediate results across PCIe, incurring heavy queuing delays and limiting system performance.

Nonetheless, the multi-stage GPU-CPU design plays an important role. Recent work shows production-scale recommendation model sizes are growing rapidly—by an order of magnitude in just three years [39]. For production-scale models that are larger than the DRAM capacity available on GPUs (e.g., $\sim$ 15GB on NVIDIA T4), designers will need to decompose models into multiple stages. Here, frontend stages run on the GPU in order to circumvent storage capacity limits and exploit data-parallelism with the larger input working set size; the backend models run on the CPU. Figure 8(top) shows that this multi-stage GPU-CPU design achieves up to $3\times$ lower latency than the multi-stage CPU-only configuration.

Takeaway 3: *By maximizing throughput at low latency, GPUs unlock higher recommendation quality.*

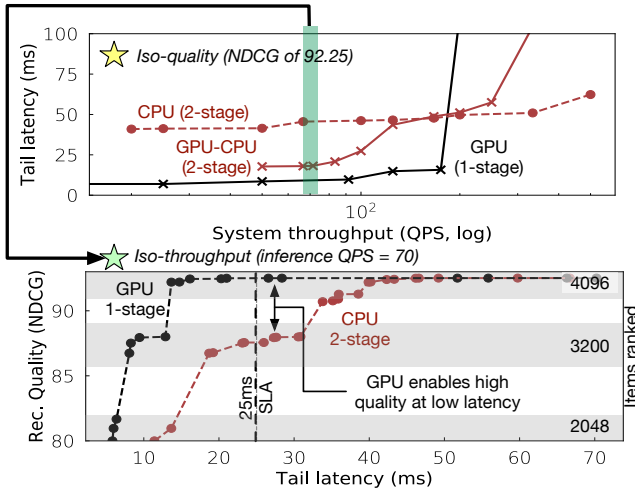Despite the GPUs achieving $3\times$ lower latency than the

6

Figure 8: (Top) At iso-quality, mapping frontend (i.e., data-parallel) stages to GPUs reduces tail-latency by up to $3\times$; CPU-only execution achieves higher system throughput. (Bottom) At a lower system throughput (i.e., QPS of 70), the lower latency on GPUs can be traded off for higher quality compared to CPU-based execution.

CPU-only designs (see Figure 8(top)), the GPUs remain underutilized with an occupancy of 25%, and memory and power utilization of 10% and 45%, respectively. Improving utilization requires higher batching. Unfortunately, as we increase batching and system throughput (x-axis), the GPU-enabled designs suffer from a sudden degradation in tail-latency due to high queuing delays; in comparison, the CPUs sustain higher throughput by concurrently processing queries across cores (e.g., task-parallelism).

While the latency reduction from GPU's does not translate to higher throughput, it can enable higher quality. Figure 8(bottom) illustrates the tradeoff between tail-latency and quality for CPU- and GPU-based recommendation at iso-throughput. Following our previous results, we show the optimal configurations: single-stage GPU-only and two-stage CPU-only designs. Given the fixed models, RecPipe tradeoffs off latency for quality by increasing the number of items ranked per query. At a strict SLA target of 25*ms*, the CPU achieves an NDCG of 87, while the GPU achieves an NDCG of 92.25. The increase in quality is a direct result of GPU's data-parallel architecture allowing it to rank 4096 items compared to the CPU ranking only 3200 items at the 25 *ms* SLA. Thus, AI accelerators for recommendation must be evaluated not only for performance benefits but also on quality achieved under strict performance and resource constraints.

Limitations of commodity hardware. Based on the performance analysis above, we identify multiple limitations of commodity platforms running multi-stage recommendation. In particular, GPUs do not directly benefit from decomposing models into multiple stages. This is due to the limits of multi-tenant execution, under utilized hardware when separately exploiting data- and model- level parallelism across stages, and high PCIe data communication between stages. Given these limitations and the growing scale of personalized recommendation across Internet services [39, 59, 60], we use RecPipe to unlock the opportunities from multi-stage ranking

by designing specialized hardware to provide high quality and infrastructure efficiency, in the following section.

6. Analysis of RecPipeAccel's Design Space

This section proposes RPAccel, a specialized accelerator tailored to multi-stage recommendation models. We start with a baseline TPU-like recommendation accelerator [26]. The baseline optimizes for low-latency single-stage inference, but suffers from low utilization and system throughput on multi-stage pipelines. To accelerate multi-stage recommendation, as summarized in Section 3.2, RPAccel comprises four main features that exploit distinct opportunities enabled by RecPipe: the pipeline execution, a reconfigurable MLP unit, a top-k filtering unit, and the partitioned embedding cache for hot-vectors across models and prefetched backend vectors.

6.1 Mapping multi-stage pipelines to RPAccel

Figure 9(left) illustrates the high-level architecture of RPAccel. Unlike prior art which accelerates *single-stage* model inferences alone, RPAccel is designed to process queries end-to-end: model inferences for *multiple stages* and *filtering* top-k user-item interactions between stages. Figure 9(center) shows how multi-stage recommendation is mapped onto RPAccel. Networks across the stages share accelerator memory and compute resources. For each stage, to produce predicted CTR scores for each user-item pair, RPAccel comprises an MLP and embedding gather unit. RPAccel implements a set of top-k filtering units to identify high-quality user-item pairs.

Takeaway 4: *Breaking queries into multiple sub-batches enables pipelined execution of frontend and backend stages.*

Figure 9(right) shows the temporal mapping of multi-stage recommendation onto RPAccel. To reduce latency, RPAccel pipelines frontend and backend stages by breaking queries into smaller sub-batches. As an example, Figure 9(right) shows RPAccel splitting a single query of 4K items into four smaller batches of 1K each, overlapping frontend and backend stages. The degree of sub-batching must be carefully balanced in order to maintain high utilization and quality. Smaller batch-sizes incur higher inference overheads (e.g., weight loading) but can better overlap frontend and backend stages. Furthermore, splitting queries into n smaller batches can degrade quality as the top-k items in each stage are set by stitching the top-$\frac{k}{n}$ items in each batch. Using RecPipe, we ensure the system maintains high-quality and splits queries into four sub-batches for workloads studied in this paper.

6.2 Customization of RPAccel micro-architecture

Below we detail RPAccel's micro-architectural design space.

Takeaway 5: *Splitting monolithic systolic arrays into sub-arrays improves recommendation inference throughput by concurrently processing multiple models and queries.*

As recommendation comprises large input working set sizes, RPAccel implements a weight stationary, systolic array-based MLP engine [6, 30, 44]. To concurrently process multiple stages and queries, RPAccel dynamically splits a monolithic array into independent sub-arrays [13]. Figure 10(a) illustrates the benefit of a reconfigurable systolic array for multi-stage recommendation. We show the MAC utilization for various array sizes and models. Larger arrays achieve lower latency but suffer from lower utilization when process-
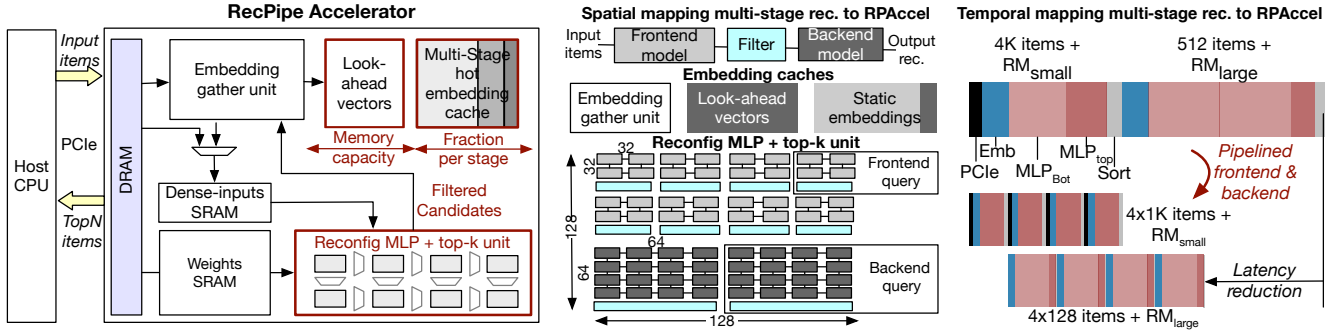
Figure 9: (Left) Overall design of the RecPipe accelerator (RPAccel) comprising an embedding gather unit with two on-chip caches for static and dynamic vectors, and a reconfigurable MLP and top-k filtering unit. (Middle) Static mapping of multi-stage recommendation onto RPAccel. Frontend and backend share both memory and compute resources. (Right) Temporal mapping of multi-stage recommendation onto RPAccel with pipelined frontend and backend models.
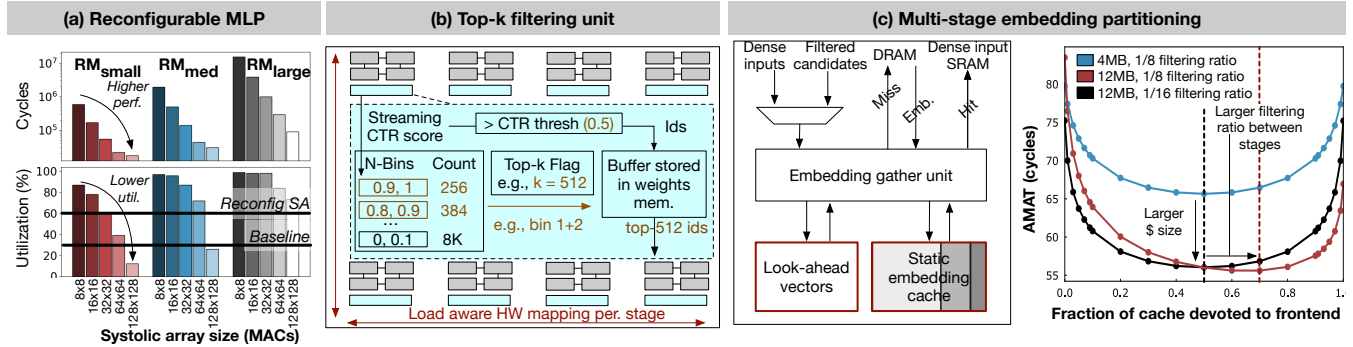

Figure 10: Design space exploration of RPAccel. (a) Larger systolic arrays suffer from low utilization on smaller models, motivating provisioning resources into sub-arrays for concurrent query processing. Compared to a monolithic array with 30% utilization, the reconfigurable array has a 60% utilization. (b) Top-k filtering unit designed to minimize area and power while eliminating host-accelerator PCIe communication overheads. (c) On-chip embedding cache resources must be asymmetrically provisioned across frontend and backend to minimize average memory access time.

ing small models (i.e., RM_{small}). In fact, when processing a two-stage pipeline, the fixed, monolithic array has an average utilization of only 30%, as it is overprovisioned for the frontend (i.e., RM_{small} ranking 4K items). Splitting the monolithic array into smaller units improves utilization to 60%, doubling throughput at comparable latency.

Note, RPAccel's reconfigurable systolic array is inspired by prior work which proposes a fission architecture to split monolithic arrays into sub-arrays for multi-tenancy [13]. Customized for multi-stage recommendation, RPAccel eliminates complex, omni-directional interconnects, incurring a lower area and power penalty (i.e., 13% area and 21% power in [13] versus 6% and 11% in RPAccel)[1], and extends the reconfigurability in response to application QPS and SLA targets.

Takeaway 6: *Implementing top-k user-item filtering units in specialized hardware eliminates PCIe overheads.*

Based on the predicted CTR, top scoring user-item interactions must be filtered and forwarded to subsequent recommendation stages. Prior recommendation accelerators only process MLP inference [26, 29]. Thus, the filtering step is offloaded to host-processors incurring high PCIe overheads [26]. To eliminate communication overheads, RPAccel implements a set of on-chip top-k filtering units (see Figure 9 middle, blue). One approach to identify the top-

---

[1]Following the baseline [13], we exclude on-chip SRAM when comparing area and power. Figure 11 includes SRAM overheads.

k user-item pairs is to sort all CTR scores. Unfortunately, sorting latency scales with the number of items to rank, potentially consuming tens-thousands of cycles for recommendation given large input sizes. Furthermore, existing hardware sorting units consume significant area and power [38].

Instead, RPAccel exploits two unique properties of recommendation inference to simplify the filtering unit. First, between stages, the final top-k user-item pairs need not be ordered—RPAccel implements an approximate, bucketing design. Second, the final MLP layer produces one CTR score per cycle leading to a streaming filtering unit design.

Figure 10(b) shows the resulting top-k filtering unit. The filtering unit maintains N bins (e.g., N=16). Each bin represents user-item pairs of a specific CTR score range between 0 and 1. As a new CTR score arrives every cycle, the filtering unit adds the user-item *id* to the corresponding bin and increments its counter. For example, Figure 10(b) shows the top bin counts user-item pairs with CTR scores between 0.9 and 1 (high quality). Based on the CTR score, the user-item pair *id* is stored in a dedicated portion of the systolic array banked weight SRAM. Storing all (4K) user-item *id* pairs consumes 12% of the weight SRAM. To reduce this overhead, RPAccel skips user-item pairs with low CTRs. Using RecPipe, we set a minimum CTR threshold of 0.5, reducing the overhead on weight memory to 3%. Once all user-item CTRs are categorized, the filtering unit identifies and copies
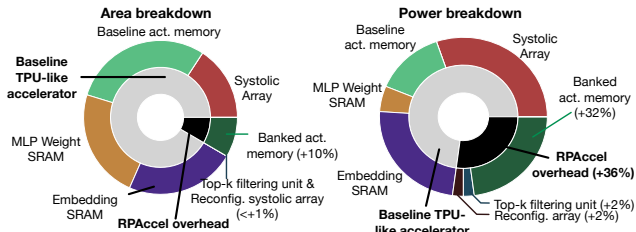
Figure 11: Compared to the baseline, RPAccel incurs 11% and 36% area (left) and power (right) overheads.

at least top-k user-item pairs indicated from the highest n bins to DRAM. These *id*s uniquely reference continuous and categorical inputs for subsequent stages.

Given the streaming design, the performance overhead of the filtering step is set by the latency it takes to identify and send user-item *id*s from the highest bins to main memory. We find this takes a couple hundred accelerator cycles, negligible compared to model inference. Although each sub-array in RPAccel's reconfigurable systolic array requires a separate top-k filtering unit, the area and power overheads are small (see Figure 11) and there is no degradation in quality.

Takeaway 7: *Asymmetrically-provisioned embedding caches tailored for each of the multi-stage models minimizes memory access latency.*

Recent work shows embedding table operations suffer from irregular memory access patterns, low compute intensity, and high storage capacities [18]. Consequently, the performance of embedding table operations is bounded by embedding vector fetch latency. Prior work exploits the power-law distribution of embedding lookups to cache frequently accessed vectors on-chip [2, 26, 33]. The embedding caches in prior work however assume a single stage recommendation model.

Instead, RPAccel implements an embedding cache customized for multi-stage recommendation by comprising (1) a *static embedding cache* that is provisioned statically for hot embedding vectors from both frontend and backend stages, (2) a *look-ahead embedding cache* that stores embedding vectors for in-flight queries. It also prefetches lookups for later stages in RPAccel's pipeline optimization (Figure 9(right)). As shown in Figure 9(left), input embedding IDs arrive either from the host processor for frontend models or the output of top-k filtering units for backend models. Based on the IDs, the embedding gather unit first checks if the corresponding vectors are in the caches. If yes, the embedding vectors are returned to the "Dense-input SRAM" to be processed by the MLP-top layers. If not, the embedding gather unit retrieves the vectors from DRAM to the look-ahead embedding cache.

Embedding cache provisioning. Following data center AI accelerators with 24MB capacity [30], we start with 16MB for embedding caches (8MB in MLP weights/activations). The size of the *look-ahead* cache is bounded by the number of items ranked in backend stages, size of embedding vectors, and maximum number of queries in flight. For the worst case we conservatively provision 4MB for the *look-ahead* cache. This leaves 12MB for the *static embedding* cache. Figure 10(c) shows the impact of asymmetrically provisioning memory for frontend and backend models on the average memory access time (AMAT) for embeddings. With a 128 byte cache line size, the embedding vector size of RM_{large},
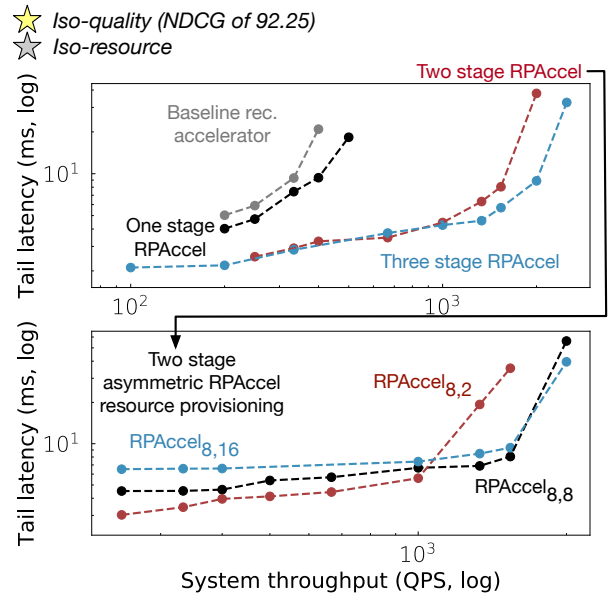

Figure 12: (Top) At iso-quality and hardware resources, co-designing multi-stage models with hardware enables lower tail-latency and higher system throughput. (Bottom) Asymmetrically provisioning RPAccel resources across stages further improves performance.

we find the fraction of storage devoted to the frontend versus backend depends on the item filtering ratio between stages. Given a filtering ratio of one-eighth for Criteo, we provision equal memory capacity for the frontend and backend.

Area and power breakdown. Figure 11 illustrates the area and power overheads of the proposed optimizations compared to the baseline, TPU-like recommendation accelerator [26]. The combination of the reconfigurable MLP unit, top-k filtering unit, and multi-stage aware embedding cache incurs a total of 11% area and 36% power overhead, moderate compared to RPAccel's performance benefits (see Section 7).

7.   Evaluation of RPAccel At-Scale

In this section we evaluate the performance of RPAccel at-scale. Instrumenting RecPipe with the simulated RPAccel we study the proposed hardware solutions in terms of quality, tail-latency, and system-throughput. We study RPAccel using publicly available models and datasets; and also project the quality and performance trends for future recommendations.

7.1   RPAccel evaluation on open-source use cases

Takeaway 8: *By accelerating multi-stage recommendation, RPAccel achieves 3× lower latency and 6× higher throughput compared to baseline, single-stage designs.*

Given fixed hardware resources, Figure 12(top) illustrates the tradeoff between throughput and latency as we vary the RPAccel-provisioning decisions for all stages. The baseline follows Centaur [26]—a single-stage recommendation accelerator which implements a TPU-like systolic array [30] and uses the host-processor to filter top-k interactions. The baseline achieves a 6*ms* and 21*ms* tail-latency at the inference throughput of 200 and 400 QPS, respectively. While achieving the same quality, the single-stage RPAccel design achieves a 4.5*ms* and 9*ms* tail-latency at 200 and 400 QPS, re-
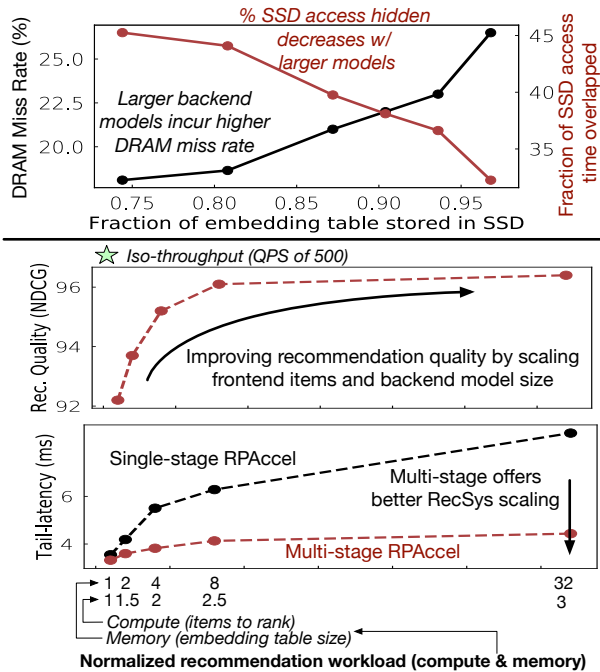
9

Figure 13: **(Top)** Projecting the performance impact of scaling recommendation models to higher capacities requiring SSD storage. **(Bottom)** Compared to the single-stage accelerator baseline, RPAccel provides graceful performance trends with future model sizes by also scaling items to rank to overlap frontend and backend stages.

spectively. Furthermore, decomposing recommendation into finer-grained pipeline enables a minimum latency of 2.1*ms* at 200 QPS or, at 6*ms* a throughput of 1300 QPS— 3× and 6× improvement over the single-stage baseline, respectively. The latency reduction and throughput increase owe RPAccel's software and hardware optimizations.

Takeaway 9: *Asymmetrically provisioning accelerator based on multi-stage recommendation models resources unlocks lower tail-latency and higher system-throughput.*

Figure 12 (bottom) illustrates the benefit of asymmetrically provisioning RPAccel resources across stages. For a two-stage recommendation pipeline, the frontend is fixed with sub-arrays while the backend implements two (i.e., RPAccel$_{8,2}$), eight (i.e., RPAccel$_{8,8}$), and sixteen sub-arrays (i.e., RPAccel$_{8,16}$). All experiments assume iso-hardware resources while achieving the maximum quality target (i.e., NDCG of 92.25).

Compared with the homogeneous accelerator (i.e., RPAccel$_{8,8}$), aggregating the backend into fewer, larger arrays RPAccel$_{8,2}$ reduces the latency at low throughput by 1.5×. Similarly, at high system load, splitting the backend into multiple, smaller units RPAccel$_{8,16}$ reduces the latency by 1.4×. Given application-level latency and system targets, asymmetrically provisioning RPAccel resources across stages widens the design space of recommendation services. Building on prior art, RPAccel resources are dynamically reconfigured to meet varying targets, given workload demands [13, 35, 48].

7.2 RPAccel evaluation on future models

So far we have analyzed the performance of RPAccel on open-source use-cases. However, recent literature shows production-scale recommendation models are rapidly growing in size, outpacing DRAM capacity and even reaching TBs in size [39]. One promising path to enabling future, production-scale models is to use higher capacity memories such as SSDs [11, 49]. Here we consider the performance implications of SSDs on RPAccel.

Storing larger embedding tables in SSD lowers embedding locality and degrades performance. Figure 13(top) shows the impact of larger embedding tables on embedding locality. While frequently accessed embedding vectors are stored DRAM, a larger portion of these tables are stored in the SSD (x-axis). For example, increasing the size of RM$_{large}$ by 32× requires storing 97% of the embedding tables in SSD. This also causes increases DRAM miss rates from 17% to 28%. Recall, RPAccel pipelines frontend and backend stages—allowing the accelerator to overlap long latency SSD accesses in the backend. However, Figure 13(top) shows with growing embedding table sizes, a smaller fraction of the accesses can be overlapped causing an increase in latency.

Takeaway 10: *Compared to baseline single-stage accelerators, RPAccel achieves higher quality and performance when scaling both frontend and backend stages towards future recommendation engines.*

In addition to scaling embedding tables in backend models (e.g., model size), one can also increase the number of items to rank in the frontend (e.g., compute demand). Figure 13(bottom) shows the impact of scaling both frontend and backend stages (x-axis) on quality. Starting from the baseline configuration, we project increasing model size by 32× and compute complexity from ranking 4K items to 12K items improves quality from an NDCG of 92.25 to 96.

Increasing the items to rank allows RPAccel to more effectively overlap the frontend and backend stages. Figure 13(bottom) shows the corresponding tail-latency impact on scaling compute and memory complexity assuming iso-throughput (QPS of 500). We show two configurations: single-stage (black) and multi-stage (red) RPAccel. By overlapping frontend and backend stages, the multi-stage design achieves higher performance for larger recommendation engines compared to the single-stage design. More generally, we show the importance of tightly-coupling algorithm and hardware scaling for future recommendation engines; RecPipe and RPAccel open such new co-design opportunities.

8. Summary of RecPipe Results

Figure 14 summarizes the performance benefits of the proposed solutions, co-designing models and hardware for multi-stage recommendation. The results show the tail-latency across three datasets (i.e., Criteo Kaggle, MovieLens 1M and 20M [20, 31]), system loads (i.e., QPS of 100, 500, 2000), and hardware platforms (i.e., CPU, GPU, Accel). The colored bars distinguish between one- (black), two- (red), and three- (blue) stage recommendation pipelines. Following our previous analysis, CPU designs assume CPU-only execution (Section 5.1). For GPU-based configurations, 1-stage designs represent GPU-only execution; 2-stage and 3-stage designs represent heterogeneous GPU-CPU execution (Section 5.2). *Accel* configurations assumes RPAccel-only execution (Section 6-7). Across the system loads and datasets, RecPipe reduces tail-latency by an average of 3.2× on com-
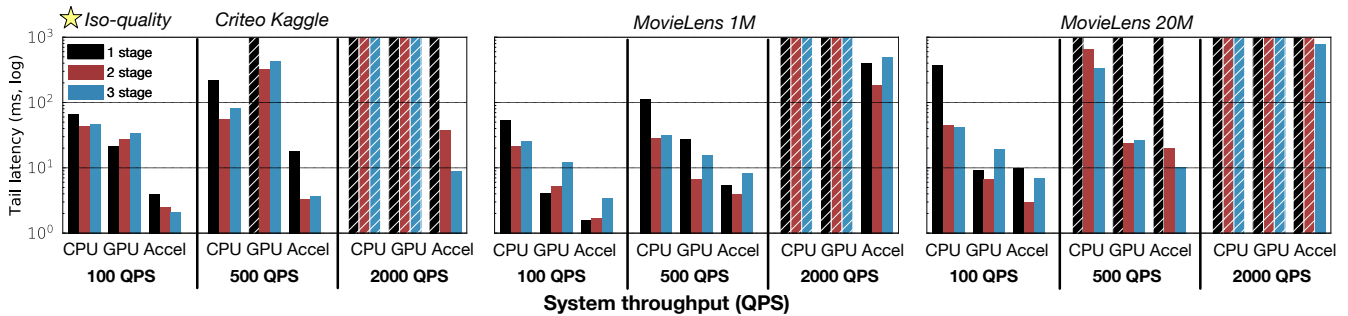
Figure 14: Summary of RecPipe results at iso-quality for the Criteo, and MovieLens 1M and 20M datasets. For each dataset, we show the tail-latency (log scale) for three system loads and hardware platforms. Configurations are greyed out when system load is not met. The optimal multi-stage design varies across loads, hardware platforms, and datasets.

modity hardware; compared to prior recommendation accelerators, RPAccel reduces tail-latency by 4.3× on average.

Differences across system loads. Across different system loads, the optimal multi-stage configuration and hardware platform varies. For instance, with the Criteo dataset on GPU-enabled hardware, between low (QPS of 100) and medium (QPS of 500) loads, the optimal number of stages varies from one to two. Similarly, for Criteo, the optimal hardware backend between low and medium loads changes from GPUs to CPUs, respectively. Differences across system loads owe to varying system optimization strategies for maximizing throughput versus minimizing latency; for example, throughput is maximized by processing multiple queries concurrently while latency is minimized by accelerating individual queries.

Differences across datasets. In addition to varying system loads, the optimal multi-stage configuration and hardware platform varies across datasets. For example with commodity hardware, on the Criteo dataset, CPUs achieve lower tail-latency than GPUs for system loads above 100 QPS; on the other hand, GPU-based designs outperform CPU-only execution for both MovieLens datasets. With RPAccel, tail-latency is optimized with the deeper three-stage pipeline for MovieLens-20M at 500 and 2000 QPS and all loads for Criteo; on MovieLens-1M however two-stage is optimal. Differences across datasets owe to the Criteo implementing DLRM [41] with higher embedding capacities while MovieLens implementing neural matrix factorization models dominated by MLP layers; furthermore, across stages the number of items to rank reduces by roughly 5×, 2.5×, and 4×, on Criteo, MovieLens 1M, and MovieLens 20M, respectively. These differences highlight the need to co-design multi-stage recommendation parameters with the underlying hardware early in the design process using frameworks like RecPipe.

Benefits of RPAccel. Compared with CPUs and GPUs, RPAccel significantly reduces tail-latency of multi-stage recommendation across different datasets and system loads. In fact, in many cases (e.g., Criteo and MovieLens20M datasets) RPAccel is optimized with deeper pipelines compared to commodity GPUs; This is a direct result of extracting data-level and model-level parallelism opportunities across multi-stage recommendation and eliminating high-overhead host-accelerator communications that RPAccel enables.

9. Related Work

While systems and computer architecture researchers have

proposed various solutions to optimize cloud-scale personalized recommendation models, relatively little work explores co-design opportunities between models and hardware to jointly optimize quality and performance, as well as the unique characteristics of multi-stage recommendation.

DNN-based recommendation models. To improve content personalization, recommendation models are growing rapidly in size and complexity [39, 58, 59, 60]. Tackling the growing model sizes, researchers have proposed techniques to compress embedding tables while preserving accuracy [12, 14, 46, 52]. Alternatively, one can decompose large monolithic models into multi-stage pipelines. Industry publications show multi-stage designs are used for serving content on Youtube [58] and Instagram [15, 27]. To balance recommendation quality and model complexity, machine learning researchers have explored a variety of modeling techniques to train each stage of the multi-stage pipeline [32]. However, in prior work, the multi-stage recommendation systems are designed to maximize quality, independent on the underlying hardware. RecPipe extends prior art by co-designing the multi-stage models and underlying hardware—commodity and specialized—in order to tightly co-optimize quality, tail-latency, and throughput for data center scale deployment.

Specialized recommendation hardware. Lots of research effort has been devoted to design specialized hardware for deep learning—especially MLPs, CNNs, and RNNs [4, 6, 7, 8, 17, 19, 30, 42, 43, 44, 47, 54, 55]. However, recommendation systems pose distinct challenges owing to their network architectures and use cases [18, 25, 41]. Given its importance, hardware proposals for accelerating recommendation models have begun to emerge [1, 3, 9, 10, 16, 26, 29, 33, 34, 36, 37, 40, 50]. While prior work focuses on improving hardware efficiency given fixed workloads, RecPipe brings quality into the mix. Accounting for both quality and performance, this work co-designs multi-stage models and hardware. In addition to RecPipe's post-training inference scheduler on commodity hardware, we compare the proposed RPAccel to a state-of-the-art TPU-like baseline recommendation accelerator, Centaur [26]; compared to the baseline, we demonstrate that by co-designing models and hardware, RPAccel jointly improves recommendation quality, tail-latency, and throughput.

10. Conclusion

Given the growing prevalence of personalized recommendation, architects have invested significant resources improv-

ing recommendation inference efficiency. While proposed solutions tackle different compute and memory bottlenecks, they do not directly co-optimize quality and performance. In this work we propose RecPipe, a system for co-designing models and hardware to jointly optimize quality, tail-latency, and throughput. First, RecPipe splits monolithic models into multi-stage pipelines exposing unique system optimization opportunities. Next, we design an inference scheduler that maps multi-stage recommendation across CPUs and GPUs. Finally, we deign a novel hardware accelerator for multi-stage recommendation which achieves high-quality while improving latency and throughput by up to $3\times$ and $6\times$, respectively, over a baseline TPU-like recommendation accelerator.

REFERENCES

[1] B. Acun, M. Murphy, X. Wang, J. Nie, C.-J. Wu, and K. Hazelwood, "Understanding training efficiency of deep learning recommendation models at scale," *arXiv preprint arXiv:2011.05497*, 2020.

[2] M. Adnan, Y. E. Maboud, D. Mahajan, and P. J. Nair, "High-performance training by exploiting hot-embeddings in recommendation systems," *arXiv preprint arXiv:2103.00686*, 2021.

[3] B. Asgari, R. Hadidi, J. Cao, D. E. Shim, S.-K. Lim, and H. Kim, "Fafnir: Accelerating sparse gathering by using efficient near-memory intelligent reduction," in *2021 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2021.

[4] A. M. Caulfield, E. S. Chung, A. Putnam, H. Angepat, J. Fowers, M. Haselman, S. Heil, M. Humphrey, P. Kaur, J.-Y. Kim *et al.*, "A cloud-scale acceleration architecture," in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2016, pp. 1–13.

[5] W. Chen, T.-y. Liu, Y. Lan, Z.-m. Ma, and H. Li, "Ranking measures and loss functions in learning to rank," in *Advances in Neural Information Processing Systems*, Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, Eds., vol. 22. Curran Associates, Inc., 2009, pp. 315–323. [Online]. Available: https://proceedings.neurips.cc/paper/2009/file/2f55707d4193dc27118a0f19a1985716-Paper.pdf

[6] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 1, pp. 127–138, 2017.

[7] Y. Chen, T. Luo, S. Liu, S. Zhang, L. He, J. Wang, L. Li, T. Chen, Z. Xu, N. Sun, and O. Teman, "Dadiannao: A machine-learning supercomputer," in *MICRO*, 2014.

[8] Y. Choi and M. Rhu, "Prema: A predictive multi-task scheduling algorithm for preemptible neural processing units," in *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2020, pp. 220–233.

[9] M. Choy, "Accelerating the modern machine learning workhorse: Recommendation inference," 2020. [Online]. Available: https://sambanova.ai/blog/accelerating-the-modern-ml-workhorse-recommendation-inference

[10] N. Corp., "Neuchips recommendation accelerator recaccel," 2020. [Online]. Available: https://2ca8d951-4386-4e41-9cab-50c86da5f5a8.filesusr.com/ugd/d79931_9382d53600f54d21a6eabe46d1f0ffa2.pdf

[11] A. Eisenman, M. Naumov, D. Gardner, M. Smelyanskiy, S. Pupyrev, K. Hazelwood, A. Cidon, and S. Katti, "Bandana: Using non-volatile memory for storing deep learning models," 2018.

[12] B. Ghaemmaghami, Z. Deng, B. Cho, L. Orshansky, A. K. Singh, M. Erez, and M. Orshansky, "Training with multi-layer embeddings for model reduction," 2020.

[13] S. Ghodrati, B. H. Ahn, J. Kyung Kim, S. Kinzer, B. R. Yatham, N. Alla, H. Sharma, M. Alian, E. Ebrahimi, N. S. Kim, C. Young, and H. Esmaeilzadeh, "Planaria: Dynamic architecture fission for spatial multi-tenant acceleration of deep neural networks," in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2020, pp. 681–697.

[14] A. Ginart, M. Naumov, D. Mudigere, J. Yang, and J. Zou, "Mixed dimension embeddings with application to memory-efficient recommendation systems," *arXiv preprint arXiv:1909.11810*, 2019.

[15] S. Goda, N. Agata, and Y. Matsumura, "A stacking ensemble model for prediction of multi-type tweet engagements," in *Proceedings of the Recommender Systems Challenge 2020*, ser. RecSysChallenge '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 6–10. [Online]. Available: https://doi.org/10.1145/3415959.3415994

[16] U. Gupta, S. Hsia, V. Saraph, X. Wang, B. Reagen, G. Wei, H. S. Lee, D. Brooks, and C. Wu, "Deeprecsys: A system for optimizing end-to-end at-scale neural recommendation inference," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, 2020, pp. 982–995.

[17] U. Gupta, B. Reagen, L. Pentecost, M. Donato, T. Tambe, A. M. Rush, G.-Y. Wei, and D. Brooks, "Masr: A modular accelerator for sparse rnns," in *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2019, pp. 1–14.

[18] U. Gupta, C.-J. Wu, X. Wang, M. Naumov, B. Reagen, D. Brooks, B. Cottel, K. Hazelwood, M. Hempstead, B. Jia *et al.*, "The architectural implications of facebook's dnn-based personalized recommendation," in *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2020, pp. 488–501.

[19] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, "EIE: Efficient inference engine on compressed deep neural network," in *Proceedings of the ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2016, pp. 243–254.

[20] F. M. Harper and J. A. Konstan, "The movielens datasets: History and context." *ACM Trans. Interact. Intell. Syst.*, vol. 5, no. 4, Dec. 2015. [Online]. Available: http://dx.doi.org/10.1145/2827872.

[21] K. Hazelwood, S. Bird, D. Brooks, S. Chintala, U. Diril, D. Dzhulgakov, M. Fawzy, B. Jia, Y. Jia, A. Kalro *et al.*, "Applied machine learning at facebook: A datacenter infrastructure perspective," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2018, pp. 620–629.

[22] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th International Conference on World Wide Web*, ser. WWW '17. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2017, pp. 173–182. [Online]. Available: https://doi.org/10.1145/3038912.3052569

[23] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 173–182.

[24] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.

[25] S. Hsia, U. Gupta, M. Wilkening, C. J. Wu, G. Y. Wei, and D. Brooks, "Cross-stack workload characterization of deep recommendation systems," in *2020 IEEE International Symposium on Workload Characterization (IISWC)*, 2020, pp. 157–168.

[26] R. Hwang, T. Kim, Y. Kwon, and M. Rhu, "Centaur: A chiplet-based, hybrid sparse-dense accelerator for personalized recommendations," in *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture*, ser. ISCA '20. IEEE Press, 2020, p. 968–981. [Online]. Available: https://doi.org/10.1109/ISCA45697.2020.00083

[27] T. G. Ivan Medvedev, Haotian Wu, "Powered by ai: Instagram's explore recommender system," 2019. [Online]. Available: https://ai.facebook.com/blog/powered-by-ai-instagrams-explore-recommender-system/

[28] K. Järvelin and J. Kekäläinen, "Cumulated gain-based evaluation of ir techniques," *ACM Transactions on Information Systems (TOIS)*, vol. 20, no. 4, pp. 422–446, 2002.

[29] W. Jiang, Z. He, S. Zhang, T. B. Preußer, K. Zeng, L. Feng, J. Zhang, T. Liu, Y. Li, J. Zhou *et al.*, "Microrec: Efficient recommendation inference by hardware and data structure solutions," *Proceedings of Machine Learning and Systems*, vol. 3, 2021.

[30] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Proceedings of the ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2017, pp. 1–12.

[31] C. Kaggle, "Display advertising challenge: Predict click-through rates on display ads," 2014. [Online]. Available: https://www.kaggle.com/c/criteo-display-ad-challenge

[32] W.-C. Kang and J. McAuley, "Candidate generation with binary codes for large-scale top-n recommendation," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 1523–1532.

[33] L. Ke, U. Gupta, B. Y. Cho, D. Brooks, V. Chandra, U. Diril, A. Firoozshahian, K. Hazelwood, B. Jia, H.-H. S. Lee *et al.*, "Recnmp: Accelerating personalized recommendation with near-memory processing," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 790–803.

[34] B. Kim, J. Park, E. Lee, M. Rhu, and J. H. Ahn, "Trim: Tensor reduction in memory," *IEEE Computer Architecture Letters*, vol. 20, no. 1, pp. 5–8, 2021.

[35] H. Kwon, L. Lai, T. Krishna, and V. Chandra, "Herald: Optimizing heterogeneous dnn accelerators for edge devices," *arXiv preprint arXiv:1909.07437*, 2019.

[36] Y. Kwon, Y. Lee, and M. Rhu, "Tensordimm: A practical near-memory processing architecture for embeddings and tensor operations in deep learning," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 740–753.

[37] Y. Kwon, Y. Lee, and M. Rhu, "Tensor casting: Co-designing algorithm-architecture for personalized recommendation training," 2020.

[38] S. Lin, P. Chen, and Y. Lin, "Hardware design of low-power high-throughput sorting unit," *IEEE Transactions on Computers*, vol. 66, no. 8, pp. 1383–1395, 2017.

[39] M. Lui, Y. Yetim, Ö. Özkan, Z. Zhao, S.-Y. Tsai, C.-J. Wu, and M. Hempstead, "Understanding capacity-driven scale-out neural recommendation inference," *arXiv preprint arXiv:2011.02084*, 2020.

[40] M. Naumov, J. Kim, D. Mudigere, S. Sridharan, X. Wang, W. Zhao, S. Yilmaz, C. Kim, H. Yuen, M. Ozdal *et al.*, "Deep learning training in facebook data centers: Design of scale-up and scale-out systems," *arXiv preprint arXiv:2003.09518*, 2020.

[41] M. Naumov, D. Mudigere, H.-J. M. Shi, J. Huang, N. Sundaraman, J. Park, X. Wang, U. Gupta, C.-J. Wu, A. G. Azzolini *et al.*, "Deep learning recommendation model for personalization and recommendation systems," *arXiv preprint arXiv:1906.00091*, 2019.

[42] L. Pentecost, M. Donato, B. Reagen, U. Gupta, S. Ma, G.-Y. Wei, and D. Brooks, "Maxnvm: Maximizing dnn storage density and inference efficiency with sparse encoding and error mitigation," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 769–781.

[43] B. Reagen, P. Whatmough, R. Adolf, S. Rama, H. Lee, S. K. Lee, J. M. Hernández-Lobato, G.-Y. Wei, and D. Brooks, "Minerva: Enabling low-power, highly-accurate deep neural network accelerators," in *Proceedings of the ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2016, pp. 267–278.

[44] A. Samajdar, Y. Zhu, P. Whatmough, M. Mattina, and T. Krishna, "Scale-sim: Systolic cnn accelerator simulator," *arXiv preprint arXiv:1811.02883*, 2018.

[45] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.

[46] H.-J. M. Shi, D. Mudigere, M. Naumov, and J. Yang, "Compositional embeddings using complementary partitions for memory-efficient recommendation systems," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 165–175.

[47] F. Silfa, G. Dot, J.-M. Arnau, and A. Gonzalez, "E-PUR: An energy-efficient processing unit for recurrent neural networks," 2017. [Online]. Available: https://arxiv.org/pdf/1711.07480.pdf

[48] L. Song, J. Mao, Y. Zhuo, X. Qian, H. Li, and Y. Chen, "Hypar: Towards hybrid parallelism for deep learning accelerator array," in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2019, pp. 56–68.

[49] M. Wilkening, U. Gupta, S. Hsia, C. Trippel, C.-J. Wu, D. Brooks, and G.-Y. Wei, "Recssd: Near data processing for solid state drive based recommendation inference," in *26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021.

[50] M. Xie, K. Ren, Y. Lu, G. Yang, Q. Xu, B. Wu, J. Lin, H. Ao, W. Xu, and J. Shu, "Kraken: Memory-efficient continual learning for large-scale real-time recommendations," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '20. IEEE Press, 2020.

[51] X. Yi, Y.-F. Chen, S. Ramesh, V. Rajashekhar, L. Hong, N. Fiedel, N. Seshadri, L. Heldt, X. Wu, and E. H. Chi, "Factorized deep retrieval and distributed tensorflow serving," ser. SysML'18, 2018.

[52] C. Yin, B. Acun, X. Liu, and C.-J. Wu, "Tt-rec: Tensor train compression for deep learning recommendation models," ser. MLSys'21, 2021.

[53] J. Zhang, S. Elnikety, S. Zarar, A. Gupta, and S. Garg, "Model-switching: Dealing with fluctuating workloads in machine-learning-as-a-service systems," in *12th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud)*, 2020.

[54] J. J. Zhang, P. Raj, S. Zarar, A. Ambardekar, and S. Garg, "Compact: On-chip compression of activations for low power systolic array based cnn acceleration," *ACM Trans. Embed. Comput. Syst.*, vol. 18, no. 5s, Oct. 2019.

[55] S. Zhang, Z. Du, L. Zhang, H. Lan, S. Liu, L. Li, Q. Guo, T. Chen, and Y. Chen, "Cambricon-X: An accelerator for sparse neural networks," in *49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Oct 2016, pp. 1–12.

[56] W. Zhao, D. Xie, R. Jia, Y. Qian, R. Ding, M. Sun, and P. Li, "Distributed hierarchical gpu parameter server for massive scale deep learning ads systems," 2020.

[57] W. Zhao, J. Zhang, D. Xie, Y. Qian, R. Jia, and P. Li, "Aibox: Ctr prediction model training on a single node," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, ser. CIKM '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 319–328.

[58] Z. Zhao, L. Hong, L. Wei, J. Chen, A. Nath, S. Andrews, A. Kumthekar, M. Sathiamoorthy, X. Yi, and E. Chi, "Recommending what video to watch next: A multitask ranking system," in *Proceedings of the 13th ACM Conference on Recommender Systems*, ser. RecSys '19. New York, NY, USA: ACM, 2019, pp. 43–51. [Online]. Available: http://doi.acm.org/10.1145/3298689.3346997

[59] G. Zhou, N. Mou, Y. Fan, Q. Pi, W. Bian, C. Zhou, X. Zhu, and K. Gai, "Deep interest evolution network for click-through rate prediction," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 5941–5948.

[60] G. Zhou, X. Zhu, C. Song, Y. Fan, H. Zhu, X. Ma, Y. Yan, J. Jin, H. Li, and K. Gai, "Deep interest network for click-through rate prediction," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2018, pp. 1059–1068.