RecNMP: Accelerating Personalized Recommendation with Near-Memory Processing

Liu Ke*, Udit Gupta, Carole-Jean Wu, Benjamin Youngjae Cho, Mark Hempstead, Brandon Reagen, Xuan Zhang*

David Brooks, Vikas Chandra, Utku Diril, Amin Firoozshahian, Kim Hazelwood, Bill Jia, Hsien-Hsin S. Lee
Meng Li, Bert Maher, Dheevatsa Mudigere, Maxim Naumov, Martin Schatz, Mikhail Smelyanskiy, Xiaodong Wang

Facebook, 1 Hacker Way, Menlo Park, CA 94025
liuke@fb.com, xuan.zhang@wustl.edu

Abstract—**Personalized recommendation systems leverage deep learning models and account for the majority of data center AI cycles. Their performance is dominated by memory-bound sparse embedding operations with unique irregular memory access patterns that pose a fundamental challenge to accelerate. This paper proposes a lightweight, commodity DRAM compliant, near-memory processing solution to accelerate personalized recommendation inference. The in-depth characterization of production-grade recommendation models shows that embedding operations with high model-, operator- and data-level parallelism lead to memory bandwidth saturation, limiting recommendation inference performance. We propose *RecNMP* which provides a scalable solution to improve system throughput, supporting a broad range of sparse embedding models. *RecNMP* is specifically tailored to production environments with heavy co-location of operators on a single server. Several hardware/software co-optimization techniques such as memory-side caching, table-aware packet scheduling, and hot entry profiling are studied, resulting in up to $9.8\times$ memory latency speedup over a highly-optimized baseline. Overall, *RecNMP* offers $4.2\times$ throughput improvement and 45.8% memory energy savings.**

I. INTRODUCTION

Personalized recommendation is a fundamental building block of many internet services used by search engines, social networks, online retail, and content streaming [7], [13], [19], [21]. Today's personalized recommendation systems leverage deep learning to maximize accuracy and deliver the best user experience [11], [22], [26], [51], [77]. The underlying deep learning models now consume the majority of the datacenter cycles spent on AI [9], [69]. For example, recent analysis reveals that the top recommendation models collectively contribute to more than 72% of all AI inference cycles across Facebook's production datacenters [9].

Despite the large computational demand and production impact, relatively little research has been conducted to optimize deep learning (DL)-based recommendation. Most research efforts within the architecture community have focused on accelerating the compute-intensive, highly-regular computational patterns found in fully-connected (FC), convolution (CNN), and recurrent (RNN) neural networks [5], [6], [10], [15], [18], [24], [27], [34], [39], [42], [48], [56], [57], [59]–[61], [63]–[68], [70], [72], [73], [75]. Unlike CNNs and RNNs,

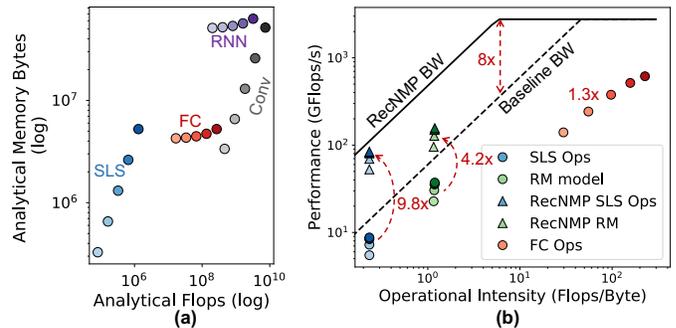*Washington University in St. Louis, work done while at Facebook.

Fig. 1. (a) Compute and memory footprint of common deep learning operators, sweeping batch size; (b) Roofline lifting effect and the operator-level (FC, SLS) and end-to-end model (RM) speedup enabled by *RecNMP*.

recommendation models exhibit low compute-intensity and little to no regularity. Existing acceleration techniques either do not apply or offer small improvements at best, as they tend to exploit regular reusable dataflow patterns and assume high spatial locality which are not the main performance bottleneck in recommendation models [69]. Given the volume of personalized inferences and their rapid growth rate occurring in the data center, an analogous effort to improve performance of these models would have substantial impact.

To suggest personalized contents to individual users, recommendation models are generally structured to take advantage of both continuous (dense) and categorical (sparse) features. The latter are captured by large embedding tables with sparse lookup and pooling operations. These embedding operations dominate the run-time of recommendation models and are markedly distinct from other layer types.

A quantitative comparison of the raw compute and memory access requirements is shown in Figure 1(a). Sparse embedding operations, represented by SparseLengthsSum (SLS), consist of a small sparse lookup into a large embedding table followed by a reduction of the embedding entries (i.e., pooling). They present two unique challenges: First, while the sparse lookup working set is comparatively small (MBs), the irregular nature of the table indices exhibits poor predictability, rendering typical prefetching and dataflow optimization techniques ineffective. Second, the embedding tables are on the order of tens to hundreds of GBs, overwhelming on-

1

chip memory resources. Furthermore, the circular points in Figure 1(b) show the operational intensity of SLS is orders of magnitude less than FC layers. Low intensity limits the potential of custom hardware including the specialized data-paths and on-chip memories used in CNN/RNN accelerators. The result is a fundamental memory bottleneck that cannot be overcome with standard caching (e.g., tiling [52]), algorithmic (e.g., input batching), or hardware acceleration techniques.

This paper proposes *RecNMP*—a near-memory processing solution to accelerate the embedding operations for DL-based recommendation. *RecNMP* is a lightweight DIMM-based system built on top of existing standard DRAM technology. We focus on DIMM-based near-memory processing [8], [23], [74] instead of resorting to specialized 2.5D/3D integration processes (e.g. HBM) [37], [39], [41]. The DIMM form factor with commodity DDR4 devices can support the 100GB+ capacities necessary for production-scale recommendation models with low cost. By eliminating the off-chip memory bottleneck and exposing higher internal bandwidth we find that *RecNMP* provides significant opportunity to improve performance and efficiency by lifting the roofline by 8× for the bandwidth-constrained region (Figure 1(b)), enabling optimization opportunity not feasible with existing systems.

We have performed a detailed characterization of recommendation models using open-source, production-scale DLRM benchmark [51], [69] as a case study. This analysis quantifies the potential benefits of near-memory processing in accelerating recommendation models and builds the intuition for co-designing the NMP hardware with the algorithmic properties of recommendation. Specifically, it highlights the opportunity for the *RecNMP* architecture in which bandwidth-intensive embedding table operations are performed in the memory and compute-intensive FC operators are performed on the CPU (or potentially on an accelerator).

The proposed *RecNMP* design exploits DIMM- and rank-level parallelism in DRAM memory systems. *RecNMP* performs local lookup and pooling functions near memory, supporting a range of sparse embedding inference operators, which produces the general Gather-Reduce execution pattern. In contrast to a general-purpose NMP architecture, we make a judicious design choice to implement selected lightweight functional units with small memory-side caches to limit the area overhead and power consumption. We combine this lightweight hardware with software optimizations including table-aware packet scheduling and hot entry profiling. Compared to previous work whose performance evaluation is solely based on randomly-generated embedding accesses [74], our characterization and experimental methodology is modeled after representative production configurations and is evaluated using real production embedding table traces. Overall, *RecNMP* leads to significant embedding access latency reduction (9.8×) and improves end-to-end recommendation inference performance (4.2×) as illustrated in Figure 1(b). Our work makes the following research contributions:

- Our in-depth workload characterization shows that production recommendation models are constrained by mem-

ory bandwidth. Our locality analysis using production embedding table traces reveals distinctive spatial and temporal reuse patterns and motivates a custom-designed NMP approach for recommendation acceleration.

- We propose *RecNMP*, a lightweight DDR4-compatible near-memory processing architecture. *RecNMP* accelerates the execution of a broad class of recommendation models and exhibits 9.8× memory latency speedup and 45.9% memory energy savings. Overall, *RecNMP* achieves 4.2× end-to-end throughput improvement.

- We examine *hardware-software co-optimization* techniques (memory-side caching, table-aware packet scheduling, and hot entry profiling) to enhance *RecNMP* performance, and customized NMP instruction with 8× DRAM command/address bandwidth expansion.

- A *production-aware evaluation framework* is developed to take into account common data-center practices and representative production configuration, such as model co-location and load balancing.

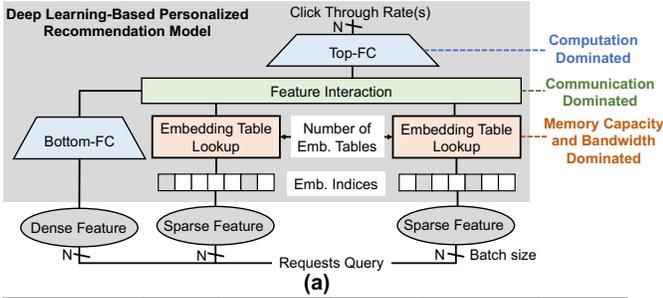## II. CHARACTERIZING DEEP LEARNING PERSONALIZED RECOMMENDATION MODELS

This section describes the general architecture of DL-based recommendation models with prominent sparse embedding features and their performance bottlenecks. As a case study, we conduct a thorough characterization of the recently-released Deep Learning Recommendation Model (DLRM) benchmark [51]. The characterization—latency breakdown, roofline analysis, bandwidth analysis, and memory locality—illustrates the unique memory requirements and access behavior of production-scale recommendation models and justifies the proposed near-memory accelerator architecture.

A. Overview of Personalized Recommendation Models

Personalized recommendation is the task of recommending content to users based on their preferences and previous interactions. For instance, video ranking (e.g., Netflix, YouTube), a small number of videos, out of potentially millions, must be recommended to each user. Thus, delivering accurate recommendations in a timely and efficient manner is important.

Most modern recommendation models have an extremely large feature set to capture a range of user behavior and preferences. These features are typically separated out into dense and sparse features. While dense features (i.e., vectors, matrices) are processed by typical DNN layers (i.e., FC, CNN, RNN), sparse features are processed by indexing large embedding tables. A general model architecture of DL-based recommendation systems is captured in Figure 2. A few examples are listed with their specific model parameters [14], [46], [51] in Figure 2(b). Similar mixture of dense and sparse features are broadly observable across many alternative recommendation models [1], [12], [14], [46], [51], [76].

Embedding table lookup and pooling operations provide an abstract representation of sparse features learned during training and are central to DL-based recommendation models. Embedding tables are organized as a set of potentially millions of vectors. Generally, embedding table operations exhibit

Fig. 2. (a) Simplified model-architecture reflecting production-scale recommendation models; (b) Parameters of representative recommendation models.

	Models	# Embs	Emb size	Pooling	Batch Size	FC Layers
DLRM	RM1-small	8	1,000,000	20-80	1-256	6
	RM1-large	12	1,000,000	20-80	1-256	6
	RM2-small	24	1,000,000	20-80	1-256	6
	RM2-large	64	1,000,000	20-80	1-256	6
	Fox	2	~ Millions	~ 50	-	4
	YouTube	2	~ Millions	~ 50	-	3

(b)

Gather-Reduce pattern; the specific element-wise reduction operation varies between models. For example, Caffe [2] comprises a family of embedding operations, prefixed by *Sparse-Lengths* (i.e., SparseLengthsWeightedSum8BitsRowwise), that perform a similar Gather-Reduce embedding operation with quantized, weighted summation. The SLS operator primitive is widely employed by other production-scale recommendation applications (e.g. YouTube [14] and Fox [46]). Our work aims to alleviate this performance bottleneck and improve system throughput by devising a novel NMP solution to offload the SLS-family embedding operations thus covering a broad class of recommendation systems.

B. A Case Study—Facebook's DLRM Benchmark

To demonstrate the advantages of near-memory processing for at-scale personalized recommendation models, we study Facebook's deep learning recommendation models (DLRMs) [51]. Dense features are initially processed by the BottomFC operators, while sparse input features are processed through the embedding table lookups. The output of these operatiors are combined and processed by TopFC producing a prediction of click-through-rate of the user-item pair.

This paper focuses on performance acceleration strategies for four recommendation models representing two canonical classes of the models, RM1 and RM2 [69]. These two model classes attribute to significant machine learning execution cycles at Facebook's production datacenter, RM1 over 30% and RM2 over 25% [9]. The parameters to configure them are shown in Figure 2(b). The notable distinguishing factor across these configurations is the number of the embedding tables. RM1 is a comparatively smaller model with few embedding tables; RM2 has tens of embedding tables.

In production environments, recommendation models employ three levels of parallelism, shown in Figure 3, to achieve high throughput under strict latency constraints [69]. Model-level parallelism grows by increasing the number of concurrent model inference (m) on a single machine, operator-level parallelism adds parallel threads (n) per model and data-level
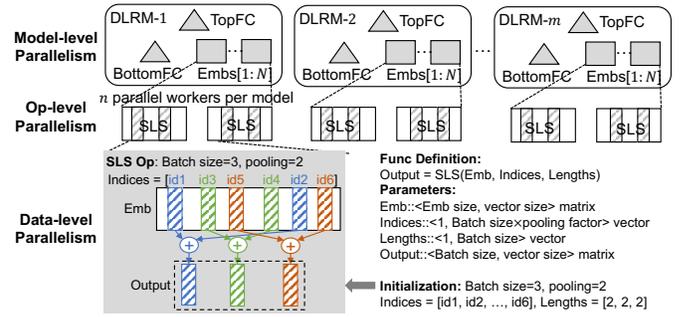

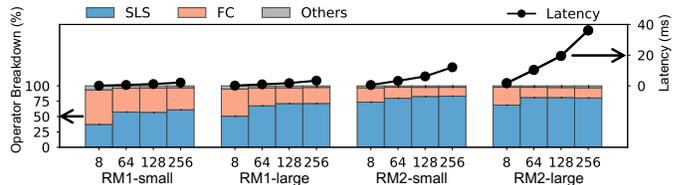Fig. 3. Model-, operator- and data-level parallelism in production system.

Fig. 4. Inference latency and breakdown across models (RM1-small, RM1-large, RM2-small, RM2-large) with varying batch sizes (8, 64, 128, 256).

parallelism is scaled by increasing batch size. An SLS operator performs a batch of pooling operations; one pooling operation performs the summation for a set of vectors. The inputs to SLS, for one batch of embedding lookups, include an indices vector containing sparse-IDs, and optionally a weight vector.

C. Operator Bottleneck Study

We observe that *the SLS-family of operators is the largest contributor to latency* in recommendation models especially as batch size, data-level parallelism, increases. Figure 4 depicts the execution time breakdown per operator with the majority of the time spent executing FC and SLS Caffe2 operators [69]. With a batch size of 8, SLS accounts for 37.2% and 50.6% of the total model execution time of RM1-small and RM1-large, respectively. Whereas for larger models represented by RM2-small and RM2-large, a more significant portion of the execution time goes into SLS (73.5%, 68.9%). Furthermore, the fraction of time spent on the embedding table operations increases with higher batch-size — 37.2% to 61.1% and 50.6% to 71.3% for RM1-small and RM1-large respectively. Note, the execution time of RM2-large is $3.6\times$ higher than RM1-large because RM2 comprises a higher number of parallel embedding tables. Generally, embedding table sizes are expected to increase further for models used in industry [74].

D. Roofline Analysis

Applying the roofline model [62], we find *recommendation models lie in the memory bandwidth-constrained region, close to the theoretical roofline performance bound*. We construct a roofline describing the theoretical limits of the test system described in Section IV. We use Intel's Memory Latency Checker (MLC)[1] to derive the memory bound. We derive the compute bound by sweeping the number of fused multiply-add (FMA) units in the processor and the operating frequency of the CPU (Turbo mode enabled).

[1]Intel MLC [31] measures the bandwidth from the processor by creating threads that traverse a large memory region in random or sequential stride as fast as possible.
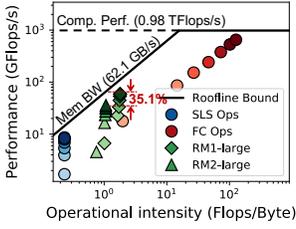
Fig. 5. Roofline of multi-threaded RM1-large and RM2-large sweeping batch size (1-256). Darker color indicates larger batch.
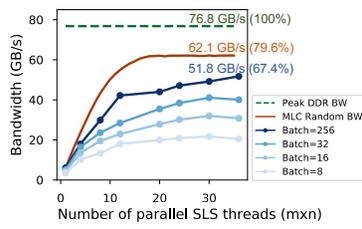
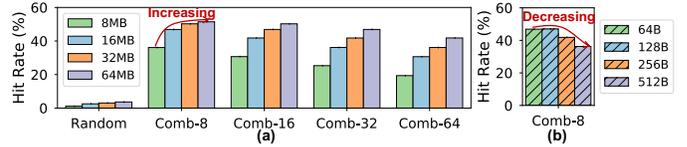Fig. 6. Memory bandwidth saturation with increasing number of parallel SLS threads and batch sizes.

Fig. 7. (a) Temporal data locality sweeping cache capacity 8-64MB with fixed cacheline size of 64B; (b) Spatial data locality sweeping cacheline size 64-512B with fixed cache capacity 16MB.

Figure 5 presents the roofline data points for the models, RM1 and RM2, as well as their corresponding FC and SLS operators separately. We sweep batch size from 1 to 256 with darker colors indicating a larger batch size. We observe that the SLS operator has low compute but higher memory requirements; the FC portion of the model has higher compute needs; and the combined model is in between. SLS has low and fixed operational intensity across batch sizes, as it performs vector lookups and element-wise summation. FC's operational intensity increases with batch size, as all requests in the batch share the same FC weights, increasing FC data reuse. With increasing batch size, the FC operator moves from the region under the memory-bound roofline to the compute-bound region. For the full model, we find RM1 and RM2 in the memory bound region, as the operational intensity is dominated by the high percentage of SLS operations. It also reveals that, with increasing batch size, the performance of SLS, as well as the RM1 and RM2 models, is approaching the theoretical performance bound of the system.

More importantly, *our roofline analysis suggests that the performance of the recommendation model is within 35.1% of the theoretical performance bound and there is little room for further improvement without increasing system memory bandwidth.* By performing the embedding lookups and pooling operations before crossing the pin-limited memory interface, near-memory processing can exploit higher internal bandwidth of the memory system, thus effectively lifting up the roofline and fundamentally improving the memory bandwidth-constrained performance bound.

E. Memory Bandwidth of Production Configurations

Executing embedding operations on real systems can saturate memory bandwidth at high model-, operator- and data-level parallelism. Figure 6 depicts the memory bandwidth consumption as we increase the number of parallel SLS threads for different batch sizes (blue curves). The green horizontal line represents the ideal peak bandwidth (76.8 GB/s, 4-channel, DDR4-2400) and the red curve is an empirical upper bound measured with Intel MLC [31]. We observe that memory bandwidth can be easily saturated by embedding operations especially as batch size and the number of threads increase. In this case, the memory bandwidth saturation point occurs (batch size = 256, thread size = 30) where more than 67.4% of the available bandwidth is taken up by SLS. In practice, a higher level of bandwidth saturation beyond

this point becomes undesirable as memory latency starts to increase significantly [36]. *What is needed is a system that can perform the Gather-Reduce operation near memory such that only the final output from the pooling returns to the CPU.*

F. Embedding Table Locality Analysis

Prior work [69] has assumed that embedding table lookups are always random, however we show, *for traces from production traffic, there exists modest level of locality mostly due to temporal reuse.* While recommendation models are limited by memory performance generally, we wanted to study the memory locality to see if caching can improve performance. We evaluate both a random trace and embedding table (T1-T8) lookup traces from production workloads used by Eisenman et al. [17]. In production systems, one recommendation model contains tens of embedding tables and multiple models are co-located on a single machine. To mimic the cache behavior of a production system, we simulate the cache hit rate for multiple embedding tables co-located on one machine. In Figure 7(a), Comb-8 means that 8 embedding tables are running on the machine and the T1-T8 traces (each for a single embedding table) are interleaved for the 8 embedding tables. For Comb-16, Comb-32 and Comb-64 we multiply the 8 embedding tables 2, 4, and 8 times on the same machine, which also approximates larger models with 16, 32 and 64 embedding tables. We use the LRU cache replacement policy and 4-way set associative cache. We assume each embedding table is stored in a contiguous logical address space and randomly mapped to free physical pages.

To estimate the amount of temporal locality present, we sweep the cache capacity between 8-64MB with fixed cacheline size of 64B. In Figure 7(a), the random trace has a low hit rate of <5% representing the worst case locality. We see that the combined simulation of production traces is much higher than random with a hit rate between 20% and 60%. More importantly, hit rate increases as cache size increases. In Section III-D, we will show how optimizations to *RecNMP* can take advantage of this locality through table-aware packet scheduling and software locality hints from batch profiling.

Spatial locality can be estimated by sweeping the cacheline size of 64-512B with a fixed cache capacity of 16MB. Figure 7(b) illustrates this sweep for the Comb-8. We observe that as the cacheline size increases, in fact, hit rate decreases. In order to isolate the effect of increased conflict misses we run the same experiment on a fully-associative cache and observe similar trends of decreasing hit rate. Thus, we conclude that embedding table lookup operations have little spatial locality.

III. RecNMP System Design

Considering the unique memory-bounded characteristics and the sparse and irregular access pattern of personalized recommendation, we propose *RecNMP*—a practical and lightweight near-memory processing solution to accelerate the dominated embedding operations. It is designed to maximize DRAM rank-level parallelism by computing directly and locally on data fetched from concurrently activated ranks.

First, we employ a minimalist style hardware architecture and embed specialized logic units and a rank-level cache to only support the SLS-family inference operators instead of general-purpose computation. The modified hardware is limited to the buffer chip within a DIMM without requiring any changes to commodity DRAM devices. Next, the sparse, irregular nature of embedding lookups exerts a high demand on command/address (C/A) bandwidth. This is addressed by sending a compressed instruction format over the standard memory interface, conforming to the standard DRAM physical pin-outs and timing constraints. Other proposed NMP solutions have employed special NMP instructions without address the C/A limitation of irregular and low spatial locality memory accesses pattern [23], [74]. We also present a hardware/software (HW/SW) interface for host-NMP coordination by adopting a heterogeneous computing programming model, similar to OpenCL [32]. Finally, we explore several HW/SW co-optimization techniques–*memory-side caching*, *table-aware scheduling* and *hot entry profiling*–that provide additional performance gains. These approaches leverage our observations from the workload characterization in the previous section.

A. Hardware Architecture

System overview. *RecNMP* resides in the buffer chip on the DIMM. The buffer chip bridges the memory channel interface from the host and the standard DRAM device interface, using data and C/A pins, as illustrated in Figure 8(a). Each buffer chip contains a RecNMP processing unit (PU) made up of a DIMM-NMP module and multiple rank-NMP modules. This approach is non-intrusive and scalable, as larger memory capacity can be provided by populating a single memory channel with multiple RecNMP-equipped DIMMs. Multiple DRR4 channels can also be utilized with software coordination.

The host-side memory controller communicates with a RecNMP PU by sending customized compressed-format NMP instructions (NMP-Inst) through the conventional memory channel interface; the PU returns the accumulated embedding pooling results (DIMM.Sum) to the host. Regular DDR4-compatible C/A and data signals (DDR.C/A and DDR.DQ) are decoded by the RecNMP PU from the NMP-Insts and then sent to all DRAM devices across all parallel ranks in a DIMM. By placing the logic at rank-level, *RecNMP* is able to issue concurrent requests to the parallel ranks and utilize, for SLS-family operators, the higher internal bandwidth present under one memory channel. Its effective bandwidth thus aggregates across all the parallel activated ranks. For example, in Figure 8(a), a memory configuration of 4 DIMMs×2 ranks per DIMM could achieve 8× higher internal bandwidth.

The DIMM-NMP module first receives a NMP-Inst through DIMM interface and then forwards it to the corresponding rank-NMP module based on the rank address. The rank-NMPs decode and execute the NMP-Inst to perform the local computation of the embedding vectors concurrently. We do not confine a SLS operation to a single rank but support aggregation across ranks within the PU. This simplifies the memory layout and increases bandwidth. DIMM-NMP performs the remaining element-wise accumulation of the partial sum vectors (PSum) from parallel ranks to arrive at the final result (DIMM.Sum). In the same fashion, Psums could be accumulated across multiple RecNMP PUs with software coordination. We will next dive into the design details on the DIMM-NMP and rank-NMP modules. While they are on the same buffer chip, having separate logical modules makes it easy to scale to DIMMs with a different number of ranks.

DIMM-NMP Module. To dispatch the NMP-Inst received from the DIMM interface, the DIMM-NMP module employs DDR PHY and protocol engine similar to the design of a conventional DIMM buffer chip relaying the DRAM C/A and DQ signals from and to the host-side memory controller. The instruction is multiplexed to the corresponding ranks based on the Rank-ID as shown in Figure 8(b). DIMM-NMP buffers the Psum vectors accumulated by each rank-NMP in its local registers and performs final summation using an adder tree before sending the final result back to the host via the standard DIMM interface. Depending on the memory system configuration, the number of ranks within a DIMM can vary, changing the number of inputs to the adder tree.

Rank-NMP Module. *RecNMP* uses the internal bandwidth on a DIMM to increase the effective bandwidth of embedding table operations, thus the majority of the logic is replicated for each rank. Three crucial functions are performed by the rank-NMP module—translating the NMP-Inst into low-level DDR C/A commands, managing *memory-side caching* and local computation of SLS-family operators. As illustrated in Figure 8(c), the NMP-Inst is decoded to control signals and register inputs. To address C/A bus limitations, all of the DDR commands for a single SLS vector is embedded in one NMP-Inst. Three fields in NMP-Inst (Figure 8(d))—DDR cmd (the presence/absence of {ACT, RD, PRE} with bit 1/0), vector size (vsize), and DRAM address (Daddr)—determine the DDR command sequence and the burst length. These are fed to the local command decoder (Rank.CmdDecoder) to generate standard DDR-style ACT/RD/PRE commands to communicate with DRAM devices. The tags are set at runtime by the host-side memory controller based on the relative physical address location of consecutive embedding accesses. This keeps the CmdDecoder in rank-NMP lightweight, as the host-side memory controller has performed the heavy-lifting tasks of request reordering, arbitration, and clock and refresh signal generation. If a 128B vector (vsize=2) requires ACT/PRE from a row buffer miss, the command sequence to DRAM devices for the NMP-Inst is {PRE, ACT Row, RD Col, RD Col+8} decoded from {ACT, RD, PRE} and vsize tags.

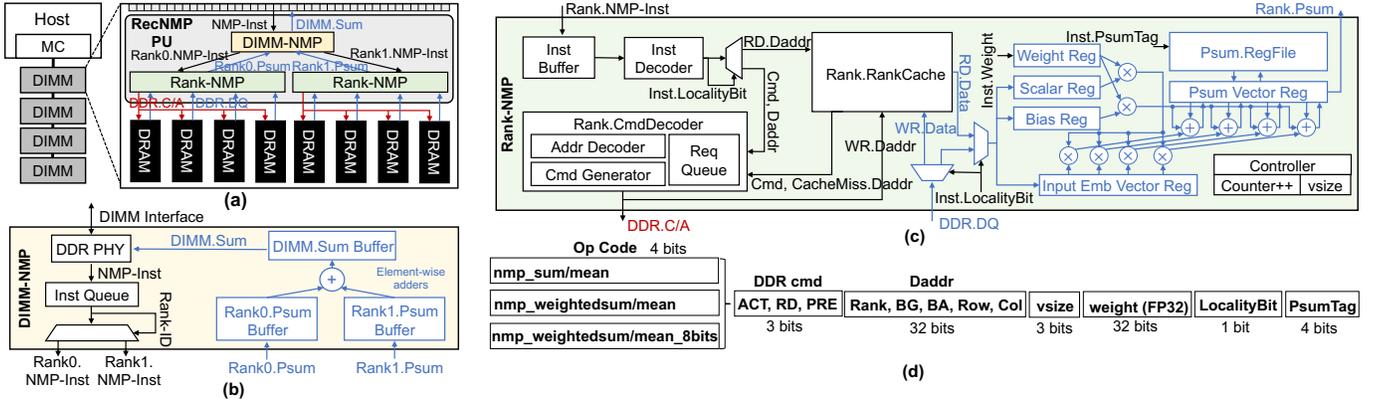Our locality analysis in Section II shows that the modest

Fig. 8. (a) Architecture overview of *RecNMP* architecture; (b) DIMM-NMP; (c) Rank-NMP; (d) NMP instruction format.

temporal locality within some embedding tables as vectors are reused. The operands of each SLS-family operator vary so caching the final result in the DIMM or CPU will be ineffective. We incorporate a memory-side cache (RankCache) in each rank-NMP module to exploit the embedding vectors reuse. The RankCache in *RecNMP* takes hints from the LocalityBit in the NMP-Inst to determine whether an embedding vector should be cached or bypassed. The detailed method to generate the LocalityBit hint through hot entry profiling will be explained in Section III-D. Entries in RankCache are tagged by the DRAM address field (Daddr). If the LocalityBit in the NMP-Inst indicates low locality, the memory request bypasses the RankCache and is forwarded to Rank.CmdDecoder to initiate a DRAM read. Embedding tables are read-only during inference, so this optimization does not impact correctness.

The datapath in the rank-NMP module supports a range of SLS-family operators. The embedding vectors returned by the RankCache or DRAM devices are loaded to the input embedding vector registers. For weighted sum computation, the weight registers are populated by the weight fields from the NMP-Inst. For quantized operators such as the SLS-8bits operator, the dequantized parameters $Scalar$ and $Bias$ are stored with the embedding vectors and can be fetched from memory to load to the Scalar and Bias registers. The Weight and Scalar/Bias registers are set to be 1 and 1/0 during execution of non-weighted and non-quantized SLS operators. The PsumTag decoded from the NMP-Inst is used to identify the embedding vectors belonging to the same pooling operations, as multiple poolings in one batch for one embedding table could be served in parallel. The controller counter, vector size register, and final sum registers in the both the DIMM-NMP and rank-NMP modules are all memory-mapped, easily accessible and configurable by the host CPU.

B. C/A Bandwidth Expansion

Although the theoretical aggregated internal bandwidth of *RecNMP* scales linearly with the number of ranks per channel, in practice, the number of concurrently activated ranks is limited by the C/A bandwidth. Due to frequent row buffer misses/conflicts from low spatial locality, accessing the embedding table entries in memory requires a large number of ACT and PRE commands. The reason is that the probability of accessing two embedding vectors in the same row is quite low,
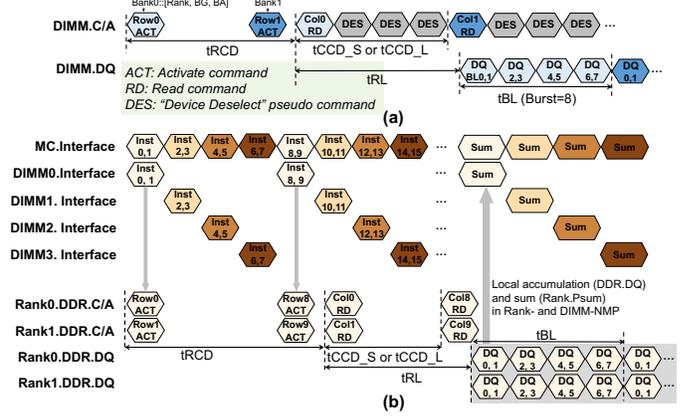

Fig. 9. Timing diagram of (a) ideal DRAM bank interleaving read operations; (b) The proposed *RecNMP* concurrent rank activation.

as spatial locality only exists in continuous DRAM data burst of one embedding vector. In production, embedding vector size ranges from 64B to 256B with low spatial locality, resulting in consecutive row buffer hits in the narrow range of 0 to 3.

To fully understand the C/A bandwidth limitation, we analyze the worst-case scenario when the embedding vector size is 64B. A typical timing diagram is presented in Figure 9(a). It shows an ideal sequence of bank-interleaved DRAM reads that could achieve one consecutive data burst. In this burst mode, the *ACT* command first sets the row address. Then the *RD* command is sent accompanied by the column address. After t_{RL} DRAM cycles, the first set of two 64-bit data (DQ0 and DQ1) appear on the data bus. The burst mode lasts for 4 DRAM cycles (burst length = 8) and transmits a total of 64B on the DQ pins at both rising and falling edges of the clock signal. Modern memory systems employ bank interleaving, therefore in the next burst cycle (4 DRAM cycles), data from a different bank can be accessed in a sequential manner. In this ideal bank interleaving case, every 64B data transfer takes 4 DRAM cycles and requires 3 DDR commands (ACT/RD/PRE) to be sent over the DIMM C/A interface, this consumes 75% of the C/A bandwidth. Activating more than one bank concurrently would require issuing more DDR commands, thus completely exhausting the available C/A bandwidth of conventional memory interface.

To overcome C/A bandwidth limitation, we propose a

customized NMP-Inst with a compressed format of DDR commands to be transmitted from memory controller to Rec-NMP PUs. Figure 9(b) illustrates the timing diagram of interleaving NMP-Inst to a 4 DIMMs × 2 Ranks per DIMM memory configuration. Eight NMP-Insts can be transferred between memory controller and DIMMs interfaces in 4 DRAM data burst cycles on double data rate. In low spatial locality case (64B embedding vector and one NMP-Inst per vector) and ideal bank interleaving, we could potentially activate 8 parallel ranks to perform 8×64B lookups concurrently in 4 DRAM data burst cycles. Although customized instructions have been proposed before [8], [23], [74], our solution is the first one to directly deal with the C/A bandwidth limitation using DDR command compression that enables up to 8× bandwidth expansion for small-sized embedding vectors (i.e. 64B) with low spatial locality. Higher expansion ratio can be achieved with larger vector size.

C. Programming Model and Execution Flow

Like previous NMP designs [23], [35], *RecNMP* adopts a heterogeneous computing programming model (e.g. OpenCL), where the application is divided into host calls running on the CPU and NMP kernels being offloaded to RecNMP PUs. NMP kernels are compiled into packets of NMP-Insts and transmitted to each memory channel over the DIMM interface to *RecNMP* PUs. Results of NMP kernels are then transmitted back to the host CPU. In Figure 8(d), each 79-bit NMP-Inst contains distinctive fields that are associated with different parameters in an embedding operation, locality hint bit (LocalityBit) and pooling tags (PsumTag) passed between the HW/SW interface. The proposed NMP-Inst format can fit within the standard 84-pin C/A and DQ interface.

Using a simple SLS function call in Figure 10(a) as an example, we walk through the execution flow of the proposed *RecNMP* programming model. First, memory is allocated for SLS input and output data, and is marked up as either Host (cacheable) or NMP (non-cacheable) regions to simplify memory coherence between the host and *RecNMP*. Variables containing host visible data, such as the two arrays *Indices and Lengths*, are initialized and loaded by the host and are cachable in the host CPU's cache hierarchy. The embedding table (Emb) in memory is initialized by the host as a host non-cacheable NMP region using a non-temporal hint (NTA) [33].

Next, the code segment marked as a NMP kernel is compiled to packets of NMP-Insts (Figure 10(b)). A single SLS NMP kernel containing one batch of embedding poolings can be split into multiple NMP packets, with each packet having one or more pooling operations. The NMP-Insts belonging to different embedding poolings in one NMP packet are tagged by PsumTag, and the maximum number of poolings in one packet is determined by the number of bits of the PsumTag. We use a 4-bit PsumTag in our design. At runtime, the NMP kernel is launched by the host with special hardware/driver support to handle NMP packet offloading; access to the memory management unit (MMU) to request memory for NMP operations; and the virtual memory system for logical-to-physical addresses translation (Figure 10(c)). The offloaded
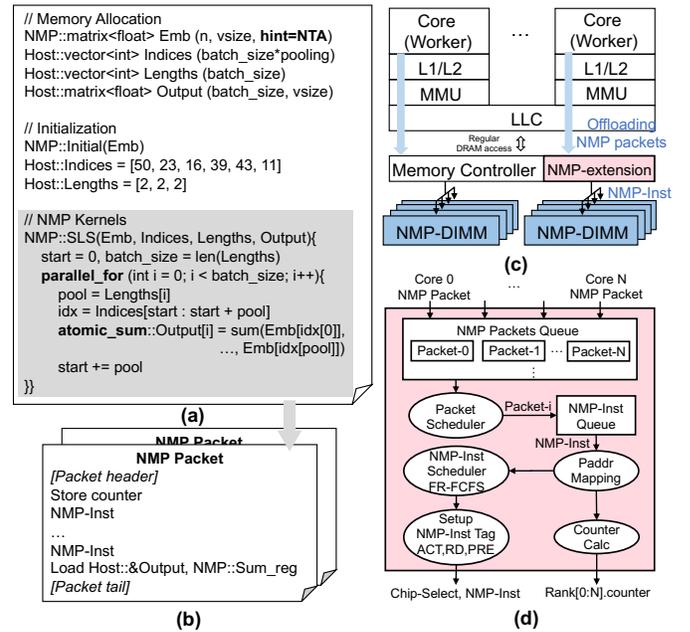

Fig. 10. (a) *RecNMP* SLS example code; (b) NMP packet; (c) NMP kernel offloading; (d) NMP-enabled memory controller.

NMP packets bypass L1/L2 and eventually arrive at the host-side memory controller with an NMP extension. To avoid scheduling the NMP packets out-of-order based on FR-FCFS policy, the NMP extension of the memory controller includes extra scheduling and arbitration logic.

As illustrated in Figure 10(d), the memory controller with the NMP extension receives concurrent NMP packets from parallel execution of multiple host cores, which are stored in a queue. Once scheduled, each NMP packet is decoded into queued NMP-Insts. Physical-to-DRAM address mapping is then performed and a FR-FCFS scheduler reorders the NMP-Insts within a packet only and not between packets. Instead of sending direct DDR commands, ACT/RD/PRE actions are compressed into the 3-bit DDR_cmd field in the NMP-Inst. The host-side memory controller also calculates the correct accumulation counter value to configure the memory-mapped control registers in the *RecNMP* PU. Finally, after the completion of all the counter-controlled local computation inside the *RecNMP* PU for one NMP packet, the final summed result is transmitted over the DIMM interface and returned to the *Output* cacheable memory region visible to the CPU.

D. HW/SW Co-optimization

Our locality analysis of production recommendation traffic in Section II-F illustrates intrinsic temporal reuse opportunities in embedding table lookups. We propose memory-side caches (RankCache) inside rank-NMP modules. To extract more performance from memory-side caching, we explore two additional HW/SW co-optimization techniques. This locality-aware optimization results in 33.7% memory latency improvement and 45.8% memory access energy saving (detailed performance benefits will be presented in Section V).

First, to preserve the intrinsic locality from embedding lookups residing in one table, we propose to prioritize scheduling NMP packets from a single batch requests to the same
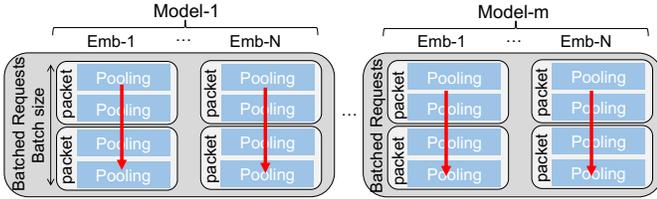
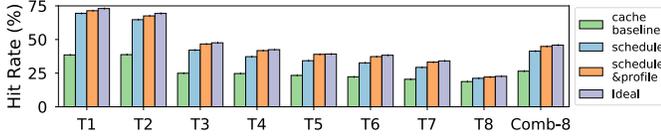Fig. 11. NMP packet scheduling scheme that prioritizes batch of single table.

Fig. 12. Hit rate of 1MB cache without optimization, with table-aware packet scheduling optimization, with both table-aware packet scheduling and hot entry profiling optimization, and ideal case without interference.

embedding table together – *table-aware packet scheduling*. In production workloads, the memory controller receives NMP packets from parallel SLS threads with equal scheduling priority. The intra-embedding table temporal locality is not easily retained because of the interference from lookup operations of multiple embedding tables. This locality can be further degraded when multiple recommendation models are co-located. Therefore, as illustrated in Figure 11, we propose an optimized table-aware NMP packet scheduling strategy to exploit the intrinsic temporal locality within a batch of requests by ordering packets from the same embedding table in one batch first, allowing the embedding vectors to be fetched together, thereby retaining the temporal locality. SLS operators access separate embedding tables as running in parallel threads, the mechanics of our implementation comes from the thread-level memory scheduler [53].

Next, we propose another optimization technique – *hot entry profiling*, built on top of the observation that a small subset of embedding entries exhibit relatively higher reuse characteristics. We profile the vector of indices used for embedding table lookup in an NMP kernel and mark the entries with high locality by explicitly annotating NMP-Insts with a *LocalityBit*. NMP-Inst with LocalityBit set will be cached in the RankCache; otherwise, the request will bypass the RankCache. This hot entry profiling step can be performed before model inference and issuing SLS requests and only costs <2% of total end-to-end execution time. We profile the indices of each incoming batch of embedding lookups and set LocalityBit if the vectors are accessed $> t$ times within the batch. Infrequent ($< t$ times) vectors will bypass the RankCache and are read directly from the DRAM devices. We sweep the threshold t and pick the value with the highest cache hit rate to use in our simulation. This hot entry profiling optimization reduces cache contention and evictions caused by the less-frequent entries in the RankCache.

Figure 12 depicts the hit rate improvement when the different optimizations are applied. Comb-8 indicates the overall hit rate at model level of 8 embedding tables (T1-T8). To gain more insights, we investigate the hit rate of embedding tables (T1 to T8) in Comb-8. The ideal bar indicates the theoretical hit rate with an infinitely sized cache. With the proposed co-
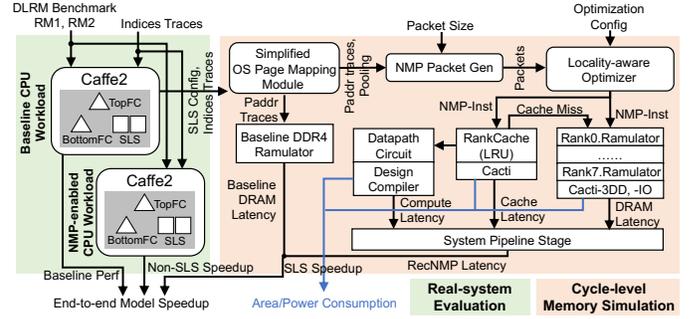

Fig. 13. *RecNMP* experimental methodology.

TABLE I
SYSTEM PARAMETERS AND CONFIGURATIONS

Real-system Configurations			
Processor	18 cores, 1.6 GHz	L1I/D	32 KB
L2 cache	1 MB	LLC	24.75 MB
DRAM	DDR4-2400MHz 8Gb ×8, 64 GB, 4 Channels × 1 DIMM × 2 Ranks, FR-FCFS 32-entry RD/WR queue, Open policy, Intel Skylake address mapping [55]		
DRAM Timing Parameters			
tRC=55, tRCD=16, tCL=16, tRP=16, tBL=4 tCCD_S=4, tCCD_L=6, tRRD_S=4, tRRD_L=6, tFAW=26			
Latency/Energy Parameters			
DDR Activate = 2.1nJ, DDR RD/WR = 14pJ/b, Off-chip IO = 22pJ/b RankCache RD/WR = 1 cycle, 50pJ/access, FP32 adder = 3 cycles, 7.89pJ/Op, FP32 mult = 4 cycles, 25.2pJ/Op			

optimization, the measured hit rate closely approaches the ideal case across the individual embedding tables, even for the trace with limited locality (T8), illustrating the proposed technique can effectively retain embedding vectors with high likelihood of reuse in RankCache.

IV. EXPERIMENTAL METHODOLOGY

Our experimental setup combines real-system evaluations with cycle-level memory simulations, as presented in Figure 13. For real-system evaluations, we run production-scale recommendation models on server-class CPUs found in the data center. This allows us to measure the impact of accelerating embedding operations as well as the side-effect of improved memory performance of FC operations on end-to-end models. Cycle-level memory simulations allow us to evaluate the design tradeoffs when DRAM systems are augmented with *RecNMP*. Table I summarizes the parameters and configurations used in the experiments. We ran experiments on an 18-core Intel Skylake with DDR4 memory. The DRAM simulation used standard DDR4 timing from a Micron datasheet [45].

Real-system evaluation. We configured the DRLM benchmark with the same model parameters and traces in Figure 2(b) and Section II. The workload characterization (Section II) and real-system experiments (Section V) are performed on single socket Intel Skylake servers, specifications in Table I.

Cycle-level memory simulation. We build the *RecNMP* cycle-level simulation framework with four main components: (1) physical addresses mapping module; (2) packet generator; (3) locality-aware optimizer; and (4) a cycle-accurate model of a RecNMP PU consisting of DRAM devices, RankCache, arithmetic and control logic. We use Ramulator [40] to conduct cycle-level evaluations of DDR4 devices.

On top of Ramulator, we build a cycle-accurate LRU cache simulator for RankCache and model of the 4-stage pipeline in the rank-NMP module. Cacti [49] is used to estimate the access latency and area/energy of RankCache. The hardware implementation used to estimate the latency, area and power of the arithmetic logic is built from Synopsys Design Compiler with a commercial 40nm technology library. To estimate the DIMM energy, we use Cacti-3DD [38] for DRAM devices and Cacti-IO [50] for off-chip I/O at the DIMM level.

During simulation we emulate the scheduling packet generation steps taken by the software stack and the memory controller. First, we apply a standard page mapping method [44] to generate the physical addresses from a trace of embedding lookups by assuming the OS randomly selects free physical pages for each logical page frame. This physical address trace is fed to Ramulator to estimate baseline memory latency. For *RecNMP* workloads, the packet generator divides the physical address trace into packets of NMP-Insts that are sent to the cycle-accurate model. Next, the when evaluating systems with HW/SW co-optimizations, the locality-aware optimizer performs table-aware packet scheduling and hot entry profiling and decides the sequence of NMP-Insts. *RecNMP* activate all memory ranks in parallel and traditional DRAM bank-interleaving is also used. For each NMP packet, performance is determined by the slowest rank that receives the heaviest memory request load. Rank-NMP and DIMM-NMP logic units are pipelined to hide the latency of memory read operations. The total latency of *RecNMP* includes extra DRAM cycles during initialization to configure the accumulation counter and the vector size register and a cycle in the final stage to transfer the sum to the host. The latency, in DRAM cycles, of the major components including RankCache, rank-NMP logic performing weighted-partial sum and final sum are in Table I.

V. EVALUATION RESULTS

This section presents a quantitative evaluation of *Rec-NMP* and shows it accelerates end-to-end personalized recommendation inference by up to $4.2\times$. We first present the latency improvement of the offloaded SLS operators on a baseline system before analysing different optimizations including placement with page coloring, memory-side caching, table-aware packet scheduling and hot-entry profiling. We compare *RecNMP* with state-of-the-art NMP systems TensorDIMM and Chameleon [23], [74]. We also analyze the effect of *RecNMP* on co-located FC operators. Finally, an end-to-end evaluation of throughput improvement and energy savings at the model level and the area/power overhead is presented.

A. SLS Operator Speedup

In theory, because *RecNMP* exploits rank-level parallelism, speedup will scale linearly with the number of ranks and number of DIMMs in a system. Therefore, we choose four memory channel configurations (# of DIMMs $\times$ # of ranks per DIMM) that correspond to 1×2, 1×4 and 2×2, and 4×2 to demonstrate a range of system implementations.

Basic *RecNMP* design without RankCache. We start by evaluating *RecNMP* without a RankCache (RecNMP-base). In
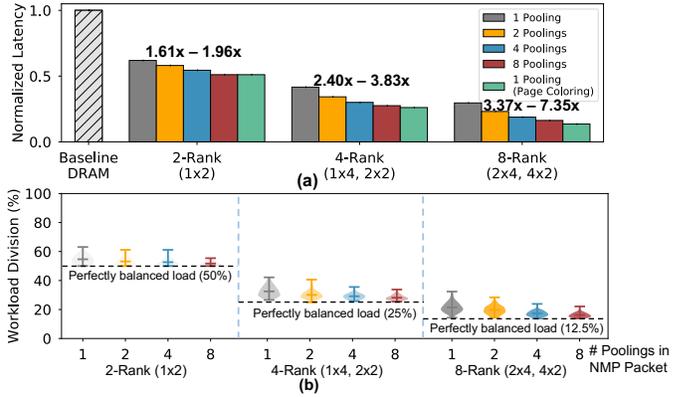

Fig. 14. (a) Normalized latency of RecNMP-base to the baseline DRAM with different memory configuration (DIMM x Rank) and NMP packet size; (b) Distribution of rank-level load imbalance for 2-, 4-, and 8-rank systems.

addition to varying the DIMM/rank configuration, we sweep the number of poolings in one NMP packet, where one pooling, in DLRM, is the sum of 80 embedding vectors. In Figure 14(a), we find 1) SLS latency indeed scales linearly as we increase the number of active ranks in a channel; 2) latency also decreases when there are more pooling operations in an NMP packet. The variation we observe, as well as the performance gap observed between the actual speedup and the theoretical speedup ($2\times$ for 2-rank, $4\times$ for 4-rank, and $8\times$ for 8-rank systems) is caused by the uneven distribution of embedding lookups across the ranks. As the ranks operate in parallel, the latency of the SLS operation is determined by the slowest rank, the rank that runs more embedding lookups. Figure 14(b) shows the statistical distribution of fraction of the work run on the slowest rank. When the NMP packet has fewer NMP-Insts, the workload distributes more unevenly, resulting in a longer tail that degrades average speedup.

To address the load imbalance, we experiment with software methods to allocate an entire embedding table to the same rank. One software approach to perform such data layout optimization is page coloring [71]. As indicated in Figure 14(a), page coloring could achieve $1.96\times$, $3.83\times$ and $7.35\times$ speedup in 2-rank, 4-rank and 8-rank system compared with the DRAM baseline. The specific page coloring mechanism can be implemented in the operating system by assigning a fixed color to the page frames used by an individual embedding table. The virtual memory system would need to be aware of the DRAM configuration to allocate pages of the same color to physical addresses that map to the same rank. This data layout optimization can lead to near-ideal speedup, but it requires maintaining high model- and task-level parallelism such that multiple NMP packets from different SLS operators can be issued simultaneously to all the available ranks.

***RecNMP* with RankCache and co-optimization.** Memory-side caching at the rank-level with table-aware packet scheduling and hot entry profiling is one of the key features of *RecNMP*; these optimizations are described in Section III-D. Figure 15(a) depicts the performance benefits (i.e. latency reduction) enabled by applying different optimization techniques: 1) adding a RankCache, 2) scheduling accesses to
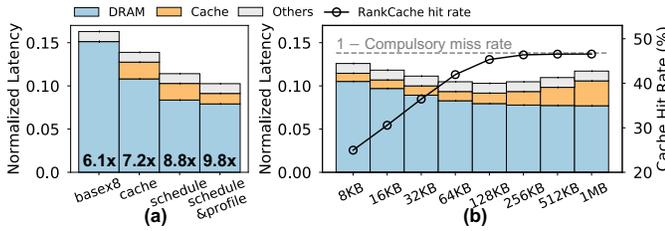
Fig. 15. (a) Normalized latency of RecNMP-cache and RecNMP-opt with schedule and hot-entry profile optimization to the baseline DRAM system; (b) Cache size sweep effects in RecNMP-opt.
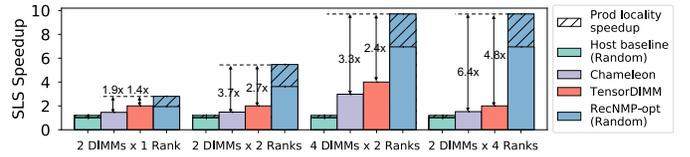

Fig. 16. Comparison between Host baseline, RecNMP-opt, TensorDIMM [74] and Chameleon [23] with both random and production traces
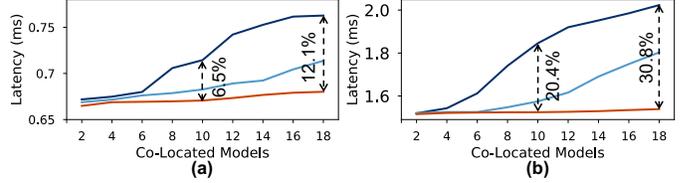

Fig. 17. Effect of model co-location on latency of (a) TopFC in RM2-small model; (b) TopFC in RM2-large model.

the same table together, 3) adding a cachability hint bit from software. Using a configuration with 8-ranks 8 poolings per packet, we observe 14.2% latency improvement by adding a 128KB RankCache and an additional 15.4% improvement by prioritizing the scheduling of NMP packets from the same table and batch. In the final combined optimization, *schedule + profile*, we pass cacheability hint after profiling the indices in the batch which reduces cache contention and allows low-locality requests not marked for caching to bypass the RankCache, delivering another 7.4% improvement. The total memory latency speedup achieved by offloading SLS to an optimized design (RecNMP-opt) is $9.8\times$.

In Figure 15(b), we sweep RankCache capacity from 8KB to 1MB and display how cache size affects the normalized latency and cache hit rate. When RankCache is small (e.g. 8KB), the low cache hit rate (e.g. 24.9%) leads to high DRAM access latency. The performance reaches the optimal design point at 128KB. Further increase of cache size has marginal improvement on hit rate, since it already reaches the compulsory limit in the trace. Yet it incurs longer cache access latency and degrades overall performance.

Performance comparison. We compare *RecNMP* with state-of-the-art NMP designs such as Chameleon [23] and TensorDIMM [74]. Both are DIMM-based near-memory processing solutions. TensorDIMM scales the embedding operation performance linearly with the number of parallel DIMMs. Since non-SLS operators are accelerated by GPUs in TensorDIMM, which is orthogonal to near-memory acceleration techniques, we only compare its memory latency speedup with *RecNMP*. Chameleon does not directly support embedding operations. We estimate its performance of Chameleon by simulating the temporal and spatial multiplexed C/A and DQ timing of Chameleon's NDA accelerators. In Figure 16, as *RecNMP* exploits rank-level parallelism, its performance scales when either the number of DIMMs and ranks increase, whereas Chameleon and TensorDIMM only scale by increasing the number of DIMMs. This is evident as we sweep the memory channel configuration. When we increase the number of ranks per-DIMM, *RecNMP* can deliver 3.3-6.4$\times$ and 2.4-4.8$\times$ better performance than Chameleon and TensorDIMM.

It is also worth noting that *RecNMP* has performance advantages (1.9$\times$ and 1.4$\times$) even in configurations with one rank per DIMM, thanks to the memory-side caching, table-aware packet scheduling, and hot-entry profiling optimization techniques. Neither Chameleon nor TensorDIMM includes a memory-side

cache to explicitly take advantage of the available locality in the memory access patterns, hence their performance, with respect to memory latency, is agnostic to traces with different amounts of data reuse. In contrast, *RecNMP* design can extract 40% more performance (shown as shaded) from production traces when compared to fully random traces.

B. FC Operator Speedup

Although *RecNMP* is designed to accelerate the execution of SLS operators, it can also improve FC performance by alleviating cache contention caused by model co-location. As the degree of data-level parallelism increases, the FC weights brought into the cache hierarchy have higher reuse, normally resulting in fewer cache misses. However, when co-located with other models, reusable FC data are often evicted early from the cache by SLS data, causing performance degradation.

Figure 17 shows the degree of performance degradation on the co-located FC operations. The amount of performance degradation experienced by the FC layers varies by the FC sizes, the degree of co-location, and the pooling values. When examining the FC performance in baseline systems, we observe worsening FC performance with larger FC weights at higher co-location degrees and higher pooling values. *RecNMP* effectively reduces the pressure from the cache contention, we show the base *RecNMP* design but RecNMP-opt impacts FC performance equally as it offloads the same SLS computation. This beneficial effect ranging from 12% to 30% is more pronounced for larger FCs whose weight parameters exceed the capacity of the L2 cache and reside mainly inside the LLC cache. For smaller FCs whose working set fits inside the L2 cache (e.g. all BottomFC and RM1's TopFC), the relative improvement is comparatively lower ($\sim 4\%$).

C. End-to-end Model Speedup

Throughput improvement. To estimate the improvement of end-to-end recommendation inference latency, we calculate the total speedup by weighting the speedup of both SLS and non-SLS operators. We measure model-level speedup across all four representative model configurations, shown in Figure 18(a). Not surprisingly, the model that spends the most time running SLS operators (RM2-large) receives the highest speedup. In Figure 18(b), the performance improvement
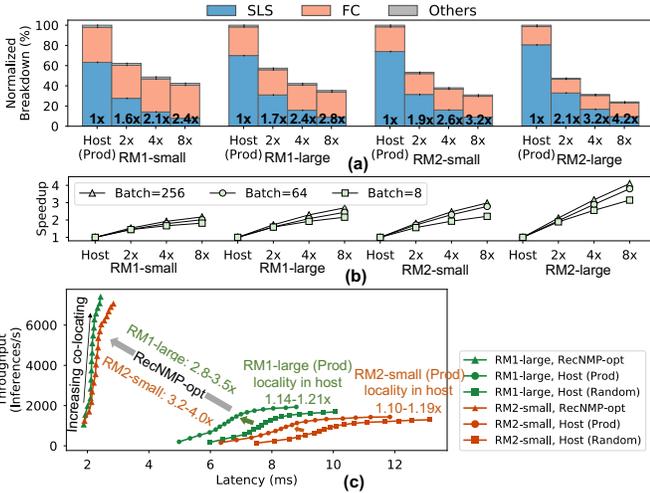
Fig. 18. (a) Single end-to-end speedup of recommendation inference with 2-rank, 4-rank and 8-rank *RecNMP* systems; (b) Single model speedup with different batch size; (c) Host and RecNMP-opt co-located model latency-throughput tradeoff.

obtained by *RecNMP* varies with batch size. In general, the model-level speedup increases with a larger batch size, as the proportion of time spent in accelerated SLS operators grows.

Figure 18(c) looks at the overall effect of increasing co-location in the presence of random or production traces for both the CPU baseline and our proposed *RecNMP* solution. Co-location generally increases the system throughput at the cost of degrading latency. Compared to random traces, the locality present in production traces improves performance. However, this locality performance "bonus" wears off as the level of model co-location increases due to the cache interference from the growing number of embedding tables in multiple models. Applying *RecNMP* in a 8-rank system results in 2.8-3.5× and 3.2-4.0× end-to-end speedup of RM1-large and RM2-small as the number of co-located models increases, because the fraction of SLS latency rises.The improvement of both latency and throughput enabled by *RecNMP* is clearly observed compared to the baseline system.

Memory energy savings. Comparing with the baseline DRAM system, *RecNMP* provide 45.8% memory energy saving. *RecNMP* saves the energy from the reduced data movement between the processor and the memory by performing local accumulation near DRAM devices and the leakage saving from reduced latency. In addition, by incorporating memory-side caching and applying co-optimization techniques to improve RankCache hit rate, *RecNMP* achieves extra energy savings by reducing the number of DRAM accesses.

Area/power overhead. We estimate *RecNMP* design overhead assuming 250MHz clock frequency and 40nm CMOS technology. The area and power numbers are derived from Synopsys Design Compiler (DC) for the arithmetic and control logic and Cacti [49] for SRAM memory (i.e. RankCache). Table II summarizes the overhead of each RecNMP processing unit for both the basic configuration without cache and the optimized configuration with cache optimization.

Compared with Chameleon, which embeds 8 CGRA cores per DIMM, our RecNMP PU consumes a fraction of the area

	RecNMP PU		Chameleon [23]
	RecNMP-base w/o RankCache	RecNMP-opt with RankCache	(8 CGRA accelerators)
Area (mm^2)	0.34	0.54	8.34
Power (mW)	151.3	184.2	3138.6-3251.8

(4.1%, 6.5% for RecNMP-base and RecNMP-opt) and power (4.6-5.9%). When scaling RecNMP PUs to multiple ranks in the DIMM, the total area and power will grow linearly, but it also translates to linearly-scaled embedding speedup. Given that a single DIMM consumes 13W [74] and a typical buffer chip takes up 100mm^2 [54], *RecNMP* incurs small area/power overhead that can easily be accommodated without requiring any changes to the DRAM devices.

VI. RELATED WORK

Performance characterization of recommendation models. Recent publications have discussed the importance and scale of personalized recommendation models in data center [1], [9], [47], [51], [69]. Compared to CNNs, RNNs, and FCs [3], [10], [24], [47], [59], the analysis demonstrates how recommendation models have unique storage, memory bandwidth, and compute requirements. For instance, [69] illustrates how Facebook's personalized recommendation models are dominated by embedding table operations. To the best of our knowledge, *RecNMP* is the first to perform locality study using production-scale models with representative embedding traces.

DRAM-based near-memory and near-data acceleration. Many prior works explore near-memory processing using 3D/2.5D-stacked DRAM technology (e.g. HMC/HBM) [4], [8], [16], [20], [28]–[30], [39], [41], [43], [58]. Due to their limited memory capacity ($16-32$GB) and high cost of ownership, these schemes are not suitable for large-scale deployment of recommendation models (10s to 100s of GBs) in production environment. Chameleon [23] introduces a practical approach to performing near-memory processing by integrating CGRA-type accelerators inside the data buffer devices in a commodity LRDIMM [23]. Unlike Chameleon's DIMM-level acceleration, *RecNMP* exploits rank-level parallelism with higher speedup potential. *RecNMP* also employs a lightweight NMP design tailored to sparse embedding operators with much lower area and power overheads than CGRAs.

System optimization for memory-constrained learning models. Sparse embedding representations have been commonly employed to augment deep neural network (DNN) models with external memory to memorize previous history. Eisenman et al. explore the use of NVMs for large embedding storage [17]. Although the proposed techniques result in $2-3×$ improvement of effective NVM read bandwidth ($2.3GB/s$), it remains far below typical DRAM bandwidth ($76.8GB/s$) and cannot fundamentally address the memory bandwidth bottleneck in recommendation models. MnnFast targets optimization for memory-augmented neural network and proposes a dedicated embedding cache to eliminate the cache contention between embedding and inference operations [25]. However, these techniques do not directly apply to

personalized recommendation consisting order-of-magnitude larger embedding tables. TensorDIMM [74] proposes a custom DIMM module enhanced with near-memory processing cores for embedding and tensor operations in deep learning. The address mapping scheme in TensorDIMM interleaves consecutive 64B within each embedding vector across the DIMM modules. Its performance thus scales at the DIMM level and relies on the inherent high spatial locality of large embedding vectors, it is unable to apply to this approach to small vectors (e.g. 64B). Given the same memory configuration, our design can outperform TensorDIMM in memory latency speedup by extracting additional performance gains from rank-level parallelism and memory-side caching optimizations. The introduction of a customized compressed NMP instruction in *RecNMP* also fundamentally addresses the C/A bandwidth constraints, without the restrictions on small embedding vectors as imposed by TensorDIMM.

VII. CONCLUSION

We propose *RecNMP*—a practical and scalable near-memory solution for personalized recommendation. We perform a systematic characterization of production-relevant recommendation models and reveal its performance bottleneck. A light-weight, commodity DRAM compliant design, *RecNMP* maximally exploits rank-level parallelism and temporal locality of production embedding traces to achieve up to $9.8\times$ performance improvement of sparse embedding operation (carried out by the SLS-family operators). Offloading SLS also offers alleviated cache contention for the non-SLS operators that remain in the CPU, resulting in up to 30% latency reduction for co-located FC operators. Overall, our system-level evaluation demonstrates that *RecNMP* offers up to $4.2\times$ throughput improvement and 45.8% memory energy saving with representative production-relevant model configurations.

REFERENCES

[1] "Breakthroughs in matching and recommendation algorithms by alibaba." [Online]. Available: https://www.alibabacloud.com/blog/breakthroughs-in-matching-and-recommendation-algorithms-by-alibaba_593976

[2] "Caffe2." [Online]. Available: https://caffe2.ai///

[3] R. Adolf, S. Rama, B. Reagen, G.-Y. Wei, and D. Brooks, "Fathom: Reference workloads for modern deep learning methods," in *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2016, pp. 1–10.

[4] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, "A scalable processing-in-memory accelerator for parallel graph processing," *ISCA*, pp. 105–117, 2015.

[5] V. Aklaghi, A. Yazdanbakhsh, K. Samadi, H. Esmaeilzadeh, and R. K. Gupta, "SnaPEA: Predictive early activation for reducing computation in deep convolutional neural networks," in *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[6] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, and A. Moshovos, "Cnvlutin: Ineffectual-neuron-free deep neural network computing," in *Proc. of the Intl. Symp. on Computer Architecture*, 2016.

[7] Amazon Personalize, https://aws.amazon.com/personalize/.

[8] Amin Farmahini-Farahani, Jung Ho Ahn, Katherine Morrow, and Nam Sung Kim, "NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules," *HPCA*, 2015.

[9] Carole-Jean Wu, David Brooks, Udit Gupta, Hsien-Hsin Lee, and Kim Hazelwood, "Deep learning: Its not all about recognizing cats and dogs," *ACM Sigarch*. [Online]. Available: https://www.sigarch.org/deep-learning-its-not-all-about-recognizing-cats-and-dogs/

[10] 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.

[11] H. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, R. Anil, Z. Haque, L. Hong, V. Jain, X. Liu, and H. Shah, "Wide & deep learning for recommender systems," in *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, DLRS RecSys 2016, Boston, MA, USA, September 15, 2016*, 2016, pp. 7–10.

[12] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir *et al.*, "Wide & deep learning for recommender systems," in *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 2016, pp. 7–10.

[13] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for youtube recommendations," in *Proceedings of the 10th ACM Conference on Recommender Systems*, ser. RecSys '16. New York, NY, USA: ACM, 2016, pp. 191–198. [Online]. Available: http://doi.acm.org/10.1145/2959100.2959190

[14] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for youtube recommendations," in *Proceedings of the 10th ACM conference on recommender systems*. ACM, 2016, pp. 191–198.

[15] C. De Sa, M. Feldman, C. Ré, and K. Olukotun, "Understanding and optimizing asynchronous low-precision stochastic gradient descent," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[16] M. Drumond, A. Daglis, N. Mirzadeh, D. Ustiugov, J. Picorel, B. Falsafi, B. Grot, and D. Pnevmatikatos, "The mondrian data engine," *ISCA*, pp. 639–651, 2017.

[17] 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," *arXiv preprint arXiv:1811.05922*, 2018.

[18] B. Feinberg, S. Wang, and E. Ipek, "Making memristive neural network accelerators reliable," in *Proc. of the Intl. Symp. on High Performance Computer Architecture*, 2018.

[19] Fortune, https://fortune.com/2019/04/30/artificial-intelligence-walmart-stores/.

[20] M. Gao, G. Ayers, and C. Kozyrakis, "Practical near-data processing for in-memory analytics frameworks," *PACT*, pp. 113–124, 2015.

[21] Google Cloud Platform, https://cloud.google.com/solutions/recommendations-using-machine-learning-on-compute-engine.

[22] H. Guo, R. Tang, Y. Ye, Z. Li, X. He, and Z. Dong, "Deepfm: An end-to-end wide & deep learning framework for CTR prediction," *CoRR*, vol. abs/1804.04950, 2018. [Online]. Available: http://arxiv.org/abs/1804.04950

[23] Hadi Asghari-Moghaddam, Young Hoon Son, Jung Ho Ahn, Nam Sung Kim, "Chameleon: Versatile and Practical Near-DRAM Acceleration Architecture for Large Memory Systems," *MICRO*, 2016.

[24] 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.

[25] Hanhwi Jang, Joonsung Kim, Jae-Eon Jo, Jaewon Lee, Jangwoo Kim, "MnnFast: a fast and scalable system architecture for memory-augmented neural networks," *ISCA*, 2019.

[26] 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 *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2018, pp. 620–629.

[27] K. Hegde, J. Yu, R. Agrawal, M. Yan, M. Pellauer, and C. W. Fletcher, "UCNN: Exploiting computational reuse in deep neural networks via weight repetition," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[28] B. Hong, G. Kim, J. H. Ahn, Y. Kwon, H. Kim, and J. Kim, "Accelerating linked-list traversal through near-data processing," *PACT*, pp. 113–124, 2016.

[29] K. Hsieh, E. Ebrahim, G. Kim, N. Chatterjee, M. O'Connor, N. Vijaykumar, O. Mutlu, and S. W. Keckler, "Transparent offloading and mapping (tom): Enabling programmer-transparent near-data processing in gpu systems," *ISCA*, pp. 204–216, 2016.

[30] K. Hsieh, S. Khan, N. Vijaykumar, K. K. Chang, A. Boroumand, S. Ghose, and O. Mutlu, "Accelerating pointer chasing in 3d-stacked memory: Challenges, mechanisms, evaluation," *ICCD*, pp. 25–32, 2016.

[31] Intel Memory Latency Checker (MLC), https://software.intel.com/en-us/articles/intelr-memory-latency-checker.

[32] J. E. Stone, D. Gohara, and G. Shi, "Opencl: A parallel programming standard for heterogeneous computing systems," *IEEE Computing in Science and Engineering*, vol. 12, no. 3, 2010.

[33] Jaekyu Lee, Hyesoon Kim, and Richard Vuduc, "When Prefetching Works, When It Doesn't, and Why," *ACM TACO*, vol. 9, no. 1, 2012.

[34] A. Jain, A. Phanishayee, J. Mars, L. Tang, and G. Pekhimenko, "Gist: Efficient data encoding for deep neural network training," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[35] Jiawen Liu, Hengyu Zhao, Matheus Almeida Ogleari, Dong Li, Jishen Zhao, "Processing-in-memory for energy-efficient neural network training: A heterogeneous approach," *MICRO*, pp. 655–668, 2018.

[36] Joseph Izraelevitz, Jian Yang, Lu Zhang, Juno Kim, Xiao Liu, Amirsaman Memaripour, Yun Joon Soh, Zixuan Wang, Yi Xu, Subramanya R. Dulloor, Jishen Zhao, Steven Swanson, "Basic Performance Measurements of the Intel Optane DC Persistent Memory Module," *arXiv preprint arXiv:1903.05714*, 2018.

[37] Junwhan Ahn, Sungjoo Yoo, Onur Mutlu, Kiyoung Choi, "Pim-enabled instructions: A low-overhead, locality-aware processing-in-memory architecture," *ISCA*, pp. 336–348, 2015.

[38] K. Chen, S. Li, N. Muralimanohar, J. H. Ahn, J. B. Brockman, and N. P. Jouppi, "Cacti-3dd: Architecture-level modeling for 3d die-stacked dram main memory," *VLSI*, pp. 33–38, 2012.

[39] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay, "Neurocube: A programmable digital neuromorphic architecture with high-density 3D memory," in *Proc. of the Intl. Symp. on Computer Architecture*, 2016.

[40] Y. Kim, W. Yang, and O. Mutlu, "Ramulator: A fast and extensible dram simulator," *IEEE Computer architecture letters*, vol. 15, no. 1, pp. 45–49, 2015.

[41] Lifeng Nai, Ramyad Hadidi, Jaewoong Sim, Hyojong Kim, Pranith Kumar, Hyesoon Kim, "Graphpim: Enabling instruction-level pim offloading in graph computing frameworks," *HPCA*, 2017.

[42] S. Liu, Z. Du, J. Tao, D. Han, T. Luo, Y. Xie, Y. Chen, and T. Chen, "Cambricon: An instruction set architecture for neural networks," in *Proc. of the Intl. Symp. on Computer Architecture*, 2016.

[43] M. Gao, J. Pu, X. Yang, M. Horowitz ,and C. Kozyrakis, "Tetris: Scalable and efficient neural network acceleration with 3d memory," *ASPLOS*, pp. 751–764, 2017.

[44] M.Gorman, "Understanding the Linux virtual memory manager." *Prentice Hall Upper Saddle River*, 2004.

[45] Micron, "MT40A2G4, MT40A1G8, MT40A512M16, 8Gb: x4, x8, x16 DDR4 SDRAM Features."

[46] Miguel Campo, Cheng-Kang Hsieh, Matt Nickens, J.J. Espinoza, Abhinav Taliyan, Julie Rieger, Jean Ho, and Bettina Sherick, "Competitive analysis system for theatrical movie releases based on movie trailer deep video representation," 2018. [Online]. Available: https://arxiv.org/abs/1807.04465

[47] MLPerf, https://mlperf.org/.

[48] N. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[49] N. Muralimanohar, R. Balasubramonian, and N. P. Jouppi, "Cacti 6.0: A tool to model large caches," *HP laboratories*, pp. 22–31, 2009.

[50] N. P. Jouppi, A. B. Kahng, N. Muralimanohar, and V. Srinivas, "Cacti-io: Cacti with off-chip power-area-timing models," *VLSI*, pp. 1254–1267, 2015.

[51] M. Naumov, D. Mudigere, H.-J. M. Shi, J. Huang, N. Sundaraman, J. Park, X. Wang, U. Gupta, C.-J. Wu, A. G. Azzolini, D. Dzhulgakov, A. Mallevich, I. Cherniavskii, Y. Lu, R. Krishnamoorthi, A. Yu, V. Kondratenko, X. Chen, V. Rao, B. Jia, L. Xiong, and M. Smelyanskiy, "Deep learning recommendation model for personalization and recommendation systems," *arXiv preprint arXiv:1906.00091*, 2019. [Online]. Available: https://arxiv.org/abs/1906.00091

[52] J. Ngiam, Z. Chen, D. Chia, P. W. Koh, Q. V. Le, and A. Y. Ng, "Tiled convolutional neural networks," *NIPS*, pp. 1279–1287, 2010.

[53] Onur Mutlu, Thomas Moscibroda, "Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors," *MICRO*, pp. 146–160, 2007.

[54] P. Meaney, L. Curley, G. Gilda, M. Hodges, D. Buerkle, R. Siegl, and R. Dong, "The IBM z13 Memory Subsystem for Big Data," *IBM Journal of Research and Development*, 2015.

[55] P. Pessl, D. Gruss, C. Maurice, M. Schwarz, and S. Mangard, "Drama: Exploiting dram addressing for cross-cpu attack," *USENIX Security Symposium*, vol. pp.565-581, 2016.

[56] A. Parashar, M. Rhu, A. Mukkara, A. Puglielli, R. Venkatesan, B. Khailany, J. Emer, S. W. Keckler, and W. J. Dally, "SCNN: An accelerator for compressed-sparse convolutional neural networks," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[57] E. Park, D. Kim, and S. Yoo, "Energy-efficient neural network accelerator based on outlier-aware low-precision computation," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[58] Q. Guo, N. Alachiotis, B. Akin, F. Sadi, G. Xu, T. M. Low, L. Pileggi, J. C. Hoe, and F. Franchetti, "3d-stacked memory-side acceleration: Accelerator and system design," *WoNDP*, 2014.

[59] 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.

[60] M. Rhu, M. O'Connor, N. Chatterjee, J. Pool, Y. Kwon, and S. W. Keckler, "Compressing DMA engine: Leveraging activation sparsity for training deep neural networks," in *Proc. of the Intl. Symp. on High Performance Computer Architecture*, 2018.

[61] M. Riera, J. M. Arnau, and A. Gonzalez, "Computation reuse in DNNs by exploiting input similarity," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[62] Samuel Williams, Andrew Waterman, and David Patterson, "Roofline: An Insightful Visual Performance Model for Floating-Point Programs and Multicore Architectures," *Communications of the ACM*, 2009.

[63] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "ISAAC: a convolutional neural network accelerator with in-situ analog arithmetic in crossbars," in *Proc. of the Intl. Symp. on Computer Architecture*, 2016, pp. 14–26.

[64] H. Sharma, J. Park, N. Suda, L. Lai, B. Chau, V. Chandra, and H. Esmaeilzadeh, "Bit fusion: Bit-level dynamically composable architecture for accelerating deep neural network," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[65] Y. Shen, M. Ferdman, and P. Milder, "Maximizing CNN accelerator efficiency through resource partitioning," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[66] M. Song, J. Zhang, H. Chen, and T. Li, "Towards efficient microarchitectural design for accelerating unsupervised GAN-based deep learning," in *Proc. of the Intl. Symp. on High Performance Computer Architecture*, 2018.

[67] M. Song, K. Zhong, J. Zhang, Y. Hu, D. Liu, W. Zhang, J. Wang, and T. Li, "In-situ AI: towards autonomous and incremental deep learning for IoT systems," in *Proc. of the Intl. Symp. on High Performance Computer Architecture*, 2018.

[68] M. Song, J. Zhao, Y. Hu, J. Zhang, and T. Li, "Prediction based execution on deep neural networks," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[69] Udit Gupta, Xiaodong Wang, Maxim Naumov, Carole-Jean Wu, Brandon Reagen, David Brooks, Bradford Cottel, Kim Hazelwood, Bill Jia, Hsien-Hsin S. Lee, Andrey Malevich, Dheevatsa Mudigere, Mikhail Smelyanskiy, Liang Xiong, Xuan Zhang, "The architectural implications of facebook's dnn-based personalized recommendation," *arXiv preprint arXiv:1906.03109*, 2019. [Online]. Available: https://arxiv.org/abs/1906.03109

[70] S. Venkataramani, A. Ranjan, S. Banerjee, D. Das, S. Avancha, A. Jagannathan, A. Durg, D. Nagaraj, B. Kaul, P. Dubey, and A. Raghunathan, "Scaledeep: A scalable compute architecture for learning and evaluating deep networks," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[71] Xiao Zhang, Sandhya Dwarkadas, Kai Shen, "Towards practical page coloring-based multicore cache management," *EuroSys*, pp. 89–102, 2009.

[72] A. Yazdanbakhsh, K. Samadi, H. Esmaeilzadeh, and N. S. Kim, "GANAX: a unified SIMD-MIMD acceleration for generative adversarial network," in *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[73] R. Yazdani, M. Riera, J.-M. Arnau, and A. Gonzalez, "The dark side of DNN pruning," *Proc. of the Intl. Symp. on Computer Architecture*, 2018.

[74] Y. L. Youngeun Kwon and M. Rhu, "Tensordimm: A practical near-memory processing architecture for embeddings and tensor operations in deep learning," *MICRO*, pp. 740–753, 2019.

[75] J. Yu, A. Lukefahr, D. J. Palframan, G. S. Dasika, R. Das, and S. A. Mahlke, "Scalpel: Customizing DNN pruning to the underlying hardware parallelism," in *Proc. of the Intl. Symp. on Computer Architecture*, 2017.

[76] 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

[77] 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, KDD 2018, London, UK, August 19-23, 2018*, 2018, pp. 1059–1068.