Tensor Manipulation Unit (TMU): Reconfigurable, Near-Memory Tensor Manipulation for High-Throughput AI SoC

Weiyu Zhou[†‡], Zheng Wang[‡*], Chao Chen[‡], Yike Li[‡§], Yongkui Yang[‡], Zhuoyu Wu[‡], Anupam Chattopadhyay[‖]

[†] Faculty of Science and Technology, University of Macau, Macau, China

[‡] Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China

[§] School of Electrical and Electronic Engineering, University College Dublin, Dublin, Ireland

[‖] College of Computing and Data Science, Nanyang Technological University, Singapore

***Abstract*—While recent advances in AI SoC design have focused heavily on accelerating tensor computation, the equally critical task of tensor manipulation—centered on high-volume data movement with minimal computation—remains underexplored. This work addresses that gap by introducing the Tensor Manipulation Unit (TMU): a reconfigurable, near-memory hardware block designed to execute data-movement-intensive (DMI) operators efficiently. TMU manipulates long datastreams in a memory-to-memory fashion using a RISC-inspired execution model and a unified addressing abstraction, enabling broad support for both coarse- and fine-grained tensor transformations. The proposed architecture integrates TMU alongside a TPU within a high-throughput AI SoC, leveraging double buffering and output forwarding to improve pipeline utilization. Fabricated in SMIC 40 nm technology, the TMU occupies only 0.019 mm^2 while supporting over 10 representative TM operators. Benchmarking shows that TMU alone achieves up to 1413.43× and 8.54× operator-level latency reduction over ARM A72 and NVIDIA Jetson TX2, respectively. When integrated with the in-house TPU, the complete system achieves a 34.6% reduction in end-to-end inference latency, demonstrating the effectiveness and scalability of reconfigurable tensor manipulation in modern AI SoCs.**

***Index Terms*—data movement, tensor manipulation, accelerator, execution model**

I. Introduction

THE previous decade has experienced the huge success of artificial neural networks (ANNs), which have been increasingly deployed on data centers, mobile edges, and terminal devices. The persistent efforts to improve peak performance and energy efficiency while reducing costs motivate the design and fabrication of customized chips for ANN. Systolic-Array-based Tensor Processing Unit (TPU) [18] is specialized in executing integer operators during inference with massive parallelism, whereas General-purpose computing on Graphics Processing Unit (GPGPU) [2] accelerates high-precision operators adopting vector-style floating-point units. Furthermore, any application inevitably contains highly flexible software code, which is conventionally deployed on a scalar-style RISC CPU. Consequently, a modern computing system or System-on-Chip (SoC) for artificial intelligence (AI) is ideally composed of scalar, vector, and tensor computing engines, plus the memory system and communication infrastructure [8].

* Corresponding author: Zheng Wang, zheng.wang@siat.ac.cn

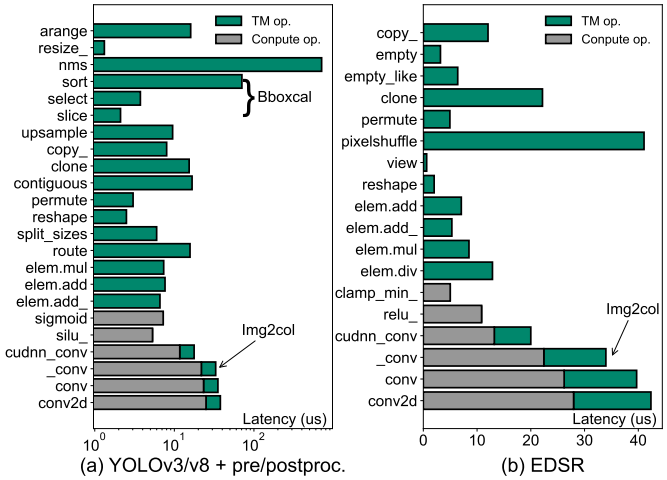

Fig. 1. The average operator-level latency when running (a) YOLOv3/v8+pre/postproc and (b) EDSR on an NVIDIA RTX 3080.

Recently, the operators have become versatile in ANNs, and their types are far beyond general matrix multiplication (GEMM), convolution, and pooling. While most of the research studies focus on accelerating the compute-intensive operators [19], very few studies and designs tend to address another rising performance bottleneck, caused by a series of data-move-intensive (DMI) operators. Existing mathematical frameworks give various names to such operators, where we name them tensor manipulation (TM), similar to array manipulation in *NumPy* [20]. Detailed in Section III, TM operators in general contain little to zero computation, but a large amount of data movement. They glue the compute-intensive operators by transforming tensors in ANN. NumPy contains 54 TM operators and the types are still growing. TM operators constitute key routines of math kernels, which are essential to pre-, post-, and intermediate processing of state-of-the-art ANN models.

As shown in Figure 1, we benchmark operator-level latency for object detection (YOLOv3/v8 [27] [10]) and super-resolution (EDSR [26]) NN models using an NVIDIA RTX 3080 GPU, which shows several TM operators are far more

costly than compute operators. operators such as Bboxcal (realized by *sort*, *select* and *slice* in GPU) and non-maximum suppression (NMS) [4] constitute the post-processing phase of most YOLO-series NNs, where *pixelshuffle* is the key operator for the EDSR model. With regard to end-to-end inference latency, TM operators account for 40.62% in EDSR when *Img2col* is included, which is used in convolution and pooling for preparing activation buffers. The root cause of such long latency is that TM operators interact intensively with off-chip memory. However, most NN accelerators move data across layers of memory hierarchy to manipulate them inefficiently.

Effectively accelerating TM operators with minimal physical resources is of increasing importance, particularly as data movement has emerged as a widely recognized performance bottleneck in modern AI workloads. However, architectural support for improving TM performance remains limited. Academic efforts such as [21] propose FPGA-based data reorganization engines for inter-memory transfers, while [22] introduces a specialized accelerator for Tucker decomposition involving tensor permutations. Other approaches [23], [24] explore in-memory data movement using 3D-stacked DRAM. Industry-wise, SoCs like NVIDIA's Jetson and Rockchip's RK3588 incorporate computer vision (CV) engines to manage tensor data, although their internal architecture details remain proprietary. Huawei's Ascend [8] includes a memory transfer engine (MTE), yet it only supports a limited subset of operators such as decomp, Img2col, and transpose.

A key limitation across prior designs is their reliance on fixed-function logic, which restricts adaptability to new and evolving TM operators. As TM grows increasingly diverse and complex, architectural reconfigurability becomes essential. Without a generalized and parameterized architectural abstraction, supporting these diverse patterns would require costly hardware redesigns or inefficient software fallbacks. This motivates the need for a reconfigurable TMU architecture that can flexibly adapt to heterogeneous TM behaviors while maintaining high throughput and low area overhead.

In this work, we propose **Tensor Manipulation Unit (TMU)** to support a series of TM operators on memory data streams. Other than a highly customized design, TMU is built under a **RISC-inspired execution model** which serves as a generic design template with reconfigurable registers to support a series of TM operators. A central feature of its design is the **Unified Address Abstraction**, which generalizes memory access patterns through parameterized address matrices. Thus operators can be encoded using shared fields with TM instructions.

TMU supports both coarse-grained and fine-grained data movements. In coarse-grained mode, the continuous datastream is uninterruptedly reorganized. In fine-grained mode, various modes of byte-level data assembling are achieved through a reconfigurable masking engine (RME). As an SoC component, TMU resides near the direct memory access (DMA) engine, which can either deliver manipulated tensor segments to tensor processing units (e.g., Img2col), or directly streaming back to off-chip DRAMs, therefore latency caused by data movement through memory hierarchies is reduced. Experiments on super-resolution and object detection NN models demonstrated a maximal speedup of 34.6% in system inference latency when TMU couples with an in-house designed TPU [11] instead of ARM CPU coupling the same TPU. Individual TM operators can be executed orders of magnitude faster on TMU compared to embedded CPU and GPU. Furthermore, TMU occupies only 0.07% silicon area of the in-house TPU, which proves its effectiveness and feasibility in modern SoCs.

The rest of the work is organized as follows. Section II surveys prior work on DMI operators. Section III introduces typical TM operators. Section IV details the design methodology for TMU. Section V illustrates the system-level integration and microarchitectural details. Section VI presents the benchmarking and synthesis results. Finally, Section VII concludes the work.

II. Related Work

In neural networks, DMI operators—such as reshape, transpose, and slice—are essential for transforming tensors across layers, optimizing memory access, and enabling efficient parallelism. Although computationally lightweight, these operators often dominate runtime due to extensive memory traffic [1].

Early efforts leveraged FPGAs to construct data reorganization engines [21], enabling coarse-grained data movement across memory hierarchies. Architectures like the Unstructured Data Processor (UDP) [7] demonstrated improved memory efficiency through dynamic data transformation. Similarly, DMA profiling and optimization techniques on FPGAs [5] reduced data transfer latency by overlapping movement with computation. Zhang et al. [22] proposed a dedicated accelerator for Tucker tensor decomposition; however, its domain specificity limits generalization to diverse DMI workloads. IBM's Active Messaging Engine (AME) [29] is a lightweight, programmable DMA core integrated into the Power10 processor. While AMEs enable efficient offloading of data movement in high-performance computing (HPC) settings, their general-purpose messaging model lacks semantic support for tensor-centric DMI operators commonly found in deep learning. ECNN [30] introduces a block-based accelerator optimized for energy-efficient CNN inference. It combines a hardware-oriented CNN model (ERNet) with a feature block instruction set (FBISA) to minimize external memory bandwidth. Although effective for fixed convolutional pipelines, its coarse execution granularity and lack of reconfigurable address control limit its applicability to dynamic TM operators such as Reshape and Bboxcal.

To reduce memory hierarchy overhead, in-situ data reorganization using 3D-stacked DRAM has been explored [23] [24]. These solutions primarily target scientific and graph-based workloads and typically require custom memory stacks. TensorCIM [3] presents a digital computing-in-memory (CIM) architecture tailored for sparse tensor operations. Its design includes modules such as the redundancy-eliminated gathering manager (REGM) and input-lookahead CIM (ILA-CIM), significantly reducing off-chip and inter-chiplet traffic. Nevertheless, its specialization for sparse and irregular data patterns limits its utility for dense and fine-grained TM operators in standard neural networks.

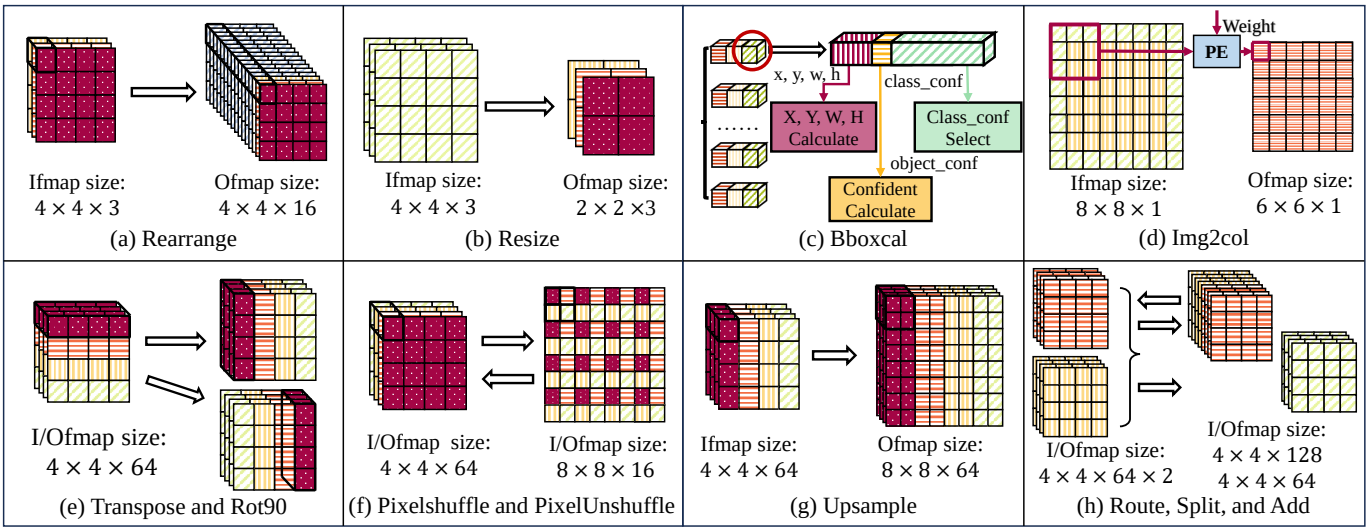

Fig. 2. Graphical representation of typical TM operators adopted in state-of-the-art neural networks.

Systolic arrays, widely used in CNN accelerators, encounter performance bottlenecks from memory bank conflicts introduced by irregular access patterns like those in Img2col. Recent approaches utilizing dynamic address generation [6] improve memory utilization, though their benefits remain largely confined to convolution-heavy workloads and do not extend to the full range of TM operators.

Commercial AI SoCs, such as NVIDIA Jetson and Rockchip RK3588, incorporate proprietary vision engines that accelerate a limited set of TM operators. Similarly, Huawei's Ascend AI processors [8] feature an MTE capable of executing operators like Img2col and transpose, but lack extensibility to support emerging operators such as Bboxcal, PixelShuffle, and Route.

Despite these advancements, existing solutions suffer from three key limitations: (1) restricted operator coverage, (2) absence of reusable architectural abstractions, and (3) weak coupling to the memory subsystem, resulting in redundant data transfers and suboptimal bandwidth utilization.

To address these challenges, the proposed TMU provides a general-purpose, near-memory acceleration framework for DMI operators. TMU employs a RISC-inspired execution model with a programmable address generation engine, enabling efficient support for a wide spectrum of coarse- and fine-grained TM operators. Operating directly on memory data streams, TMU minimizes movement overhead. Integrated alongside a TPU within a heterogeneous SoC, TMU achieves up to 34.6% system-level inference latency reduction while occupying only 0.07% of TPU area, offering a scalable and reconfigurable solution for DMI optimization in modern AI workloads.

III. Tensor Manipulation Operator

TM operators are crucial for manipulating data structures in neural networks and scientific computing. Some TM operators involve computing routines, while others are purely data arrangements. The key TM operators that have been widely adopted in state-of-the-art neural networks and implemented in the proposed TMU are highlighted as follows.

A. Fine-grained TM operator

Fine-grained TM operators function at byte-level of granularity. They are crucial for optimizing memory access patterns and data layout in DNN pre-processing or post-processing, thereby improving throughput and efficiency. Key fine-grained operators include:

1) Rearrange: Shown in Fig. 2(a), Rearrange transforms RGB data streams into higher-channel feature maps (fmaps) (e.g., 16 channels) to favor AXI burst size and DRAM access patterns. It is crucial in the preprocessing stage of vision-based NN models, where byte-level data are *fine-grained* rearranged.

2) Resize: As shown in Fig. 2(b), the bilinear interpolation operator calculates weighted averages of neighboring pixels to enable smooth image scaling. It is essential in vision AI models, ensuring visual quality with sub-pixel precision.

3) Bboxcal: As seen in Fig. 2(c), bounding boxes (Bboxes) with high confidence are extracted from YOLO's output tensors. Bboxcal is not typically accelerated by TPUs but significantly affects system inference latency. It operates on the byte level, making it a *fine-grained* operator.

4) Img2col: Critical for speeding up variants of convolution and pooling, Img2col in Fig. 2(d) extracts necessary activations from input feature maps for the activation buffers in TPU.

B. Coarse-grained TM operator

Coarse-grained TM operators manage tensors or sub-tensors whose size exceeds the hardware-defined bus width (e.g., 16 bytes for a 128-bit AXI bus), emphasizing structural or dimensional transformations. They often reshape, reorient, or combine entire feature maps to meet the architectural requirements of DNN or to fuse information from different network paths. Important coarse-grained operators include:

1) Transpose and Rot90: Transpose operators rearrange the dimensions of tensors to adapt to specific tasks, while Rot90 rotates images by 90 degrees enhancing feature representation. These operators are widely used in DNNs for managing multidimensional data, as shown in Fig. 2(e) and (f).

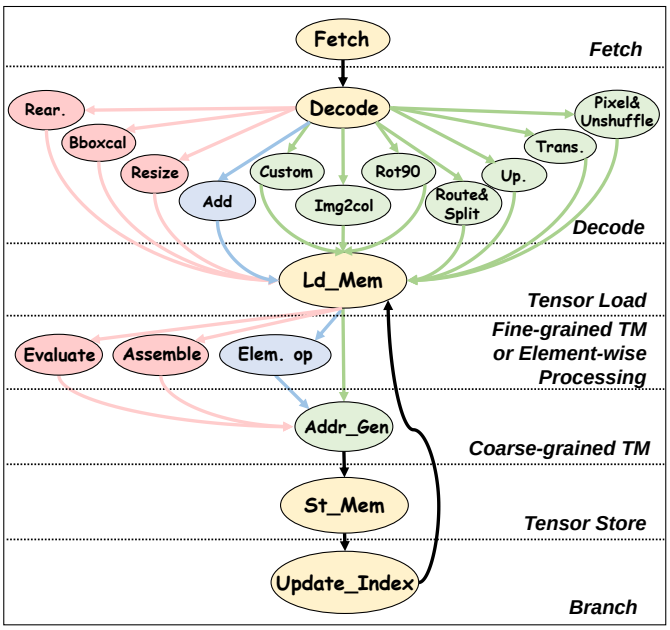

Fig. 3. Generic execution model for TM

2) PixelShuffle and PixelUnshuffle: PixelShuffle, shown in Fig. 2(g), rearranges feature maps for super-resolution, increasing width and height while reducing depth. Conversely, PixelUnshuffle, depicted in Fig. 2(h), reduces map dimensions as one of the downsampling operators.

3) Upsample: Illustrated in Fig. 2(i), upsample scales up feature maps, maintaining depth while expanding width and height, pivotal in object detection and segmentation.

4) Route, Split, and Add: Route (also known as Concat) combines feature maps along the channel dimension as depicted in Fig. 2(j). Split divides feature maps along the channel dimension as shown in Fig. 2(k). Add (also known as residual layer) performs element-wise additions with two tensors, which is essential for residual networks as depicted in Fig. 2(l).

IV. Tensor Manipulation Methodology

To raise the abstraction level of the design, we have introduced a generic tensor manipulation methodology including the execution model and address generation. The growing number of TM operators can fit into the design template following the approach outlined in this section.

A. Generic Execution Model

We propose a RISC-inspired execution model for TMU that abstracts the tensor manipulation process into eight configurable stages, as illustrated in Fig. 3. Each stage represents a distinct class of dataflow transformation or control behavior and can be selectively activated based on the characteristics of a given TM operator. This stage-based abstraction allows diverse tensor manipulations to be expressed uniformly within a unified operational framework. The functionalities of each stage are described as follows.

1) Fetch: TMU acquires instructions from local storage.

2) Decode: The instruction is issued to determine its functionality.

3) Tensor Load: TMU initiates the loading of required tensors from memory, typically from off-chip DRAM.

4) Fine-grained TM: Certain TM operators—such as Resize, Rearrange, and Bboxcal—manipulate data at the *byte-level*, often requiring flexible selection and reorganization of fine-grained tensor elements. These operations can be categorized into two high-level modes: *assemble*, which gathers selected bytes and packs them into a continuous output stream; and *evaluate*, which filters bytes based on simple comparison or thresholding and forwards only those of interest.

5) Element-wise Processing: This category includes TM operators such as Add and Mul, which perform element-wise computations across corresponding tensor positions. Each tensor element is processed independently, making these operations highly parallelizable and well-suited for fusing with adjacent computation stages.

6) Coarse-grained TM: The remaining TM operators, such as Transpose, PixelShuffle, and Split, operate on entire tensor blocks and perform structured layout transformations defined by tensor strides, shapes, and dimensions. Their access behavior can be uniformly described using a pattern-driven addressing abstraction, which encodes these transformations through a shared matrix-based formulation. This abstraction forms the foundation of the address generator discussed in the following section IV-B.

7) Tensor Store: Manipulated tensors are written back to memory (e.g., off-chip DRAM) or forwarded to downstream computing engines (e.g., TPU) for subsequent processing.

8) Branch: Long tensors cannot be manipulated in a single run. This stage updates the address and loads the following tensor segment, before proceeding to the next instruction.

Taking advantage of the above TM execution model, designers can implement various types of TM operators and enhance the versatility of the TMU. The design template also facilitates resource sharing, where there is a large similarity between TM operators. It is important to note that although our model follows a RISC-inspired pipeline structure, it is not related to or compatible with standard RISC instruction sets.

B. Unified Address Abstraction for Tensor Manipulation

A key feature of the proposed TMU is its ability to perform tensor manipulation directly in a memory-to-memory fashion, eliminating the need for CPU-driven address computations. To support a wide range of coarse-grained TM operators, the TMU employs a unified address abstraction that models memory access patterns through a parameterized and reusable formulation.

In this abstraction, each operator's access pattern is represented as an affine transformation from input tensor indices to output memory locations. Rather than implementing dedicated address generation logic for each operator, the TMU employs a shared matrix-based approach that captures stride, padding, scaling, and layout reordering in a common structure.

The unified address abstraction dynamically computes source and destination addresses at runtime, guided by

TABLE I
PARAMETERS USED IN EQ. 1 AND TABLE II.

Abbr.	Meaning	Abbr.	Meaning
$addr_{out}$	Access address	w_i	Ifmap Width
$addr_{base}$	Base address	x_k	Fmap Kernel Width
x_i	Ifmap X Position	y_k	Fmap Kernel Height
y_i	Ifmap Y Position	x_p	Fmap Padding Width
c_i	Ifmap C Position	y_p	Fmap Padding Height
x_o	Ofmap X Position	x_s	Fmap Stride Width
y_o	Ofmap Y Position	y_s	Fmap Stride Height
c_o	Ofmap C Position	s	Fmap Scale Factor

operator-specific parameters provided via the instruction stream. These parameters instantiate transformation matrices that formalize index mappings, as shown in Eq. 1, with associated terms defined in Table I.

$$addr_{out} = addr_{base} + y_o \times c_o + x_o \times c_o$$
$$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \quad (1)$$

Each TM operator corresponds to a specific pair of transformation matrices $(\mathbf{A}, \mathbf{B})$, which encode the linear relationship between input and output index triplets. For three-dimensional feature maps, the elements of **A** and **B** are typically constants, while the output indices (x_o, y_o, c_o) are computed as functions of the input indices (x_i, y_i, c_i), adapted to the semantics of each TM operator. Table II lists representative configurations of **A** and **B** for commonly used coarse-grained operators, including Transpose, Img2col, and PixelUnshuffle. This matrix-based abstraction enables flexible support for strided and dilated access, channel fusion or splitting, and other compound tensor transformations.

By encoding transformation parameters directly into TMU instruction fields, this scheme allows runtime interpretation and execution without hardware modification. As a result, the same address generation datapath can support a wide range of tensor manipulation behaviors through lightweight reconfiguration.

TABLE II
THE A AND B MATRICES OF ADDRESS GENERATION FOR COARSE-GRAINED TM OPERATORS.

TM Op.	Eq.
Transpose	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ w_i & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$
Rot90	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 0 & -1 & 0 \\ w_i & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix} + \begin{pmatrix} w_i \\ 0 \\ 0 \end{pmatrix}$
Img2col	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} \frac{1}{x_s} & 0 & 0 \\ 0 & \frac{w_i}{y_s} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix} + \begin{pmatrix} \frac{2\times x_p - x_k}{x_s} + 1 \\ \frac{2\times y_p - y_k}{y_s} + 1 \\ 0 \end{pmatrix}$
PixelShuffle	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & s \times w_i & 0 \\ 0 & 0 & \frac{1}{s} \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$
PixelUnshuffle	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} s & 0 & 0 \\ 0 & w_i & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$
Upsample	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} s & 0 & 0 \\ 0 & s \times s \times w_i & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$
Route	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & w_i & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_{i1} \\ c_{i2} \end{pmatrix}$
Split	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & w_i & 0 \\ 0 & 0 & \frac{1}{s} \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$
Add	$\begin{pmatrix} x_o \\ y_o \\ c_o \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & w_i & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ c_i \end{pmatrix}$

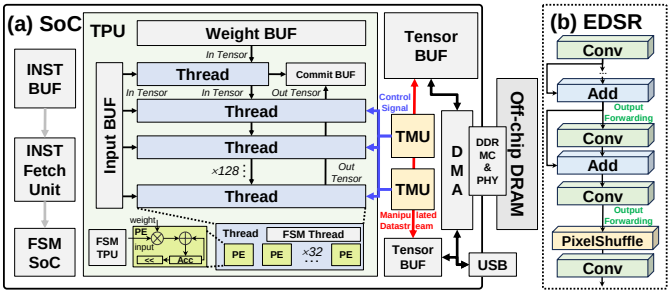

Fig. 4. (a) System architecture integrating the proposed TMU and TPU. Two TMUs are deployed to support tensor prefetching. (b) Network architecture of EDSR.

V. TMU ARCHITECTURE

This section presents the implementation of the TMU, covering its system-level integration, microarchitecture, and dataflow organization.

A. System Architecture

As illustrated in Fig. 4, the TMU serves as a key system-level component within our AI-oriented SoC architecture. The SoC executes neural networks as a sequence of operators, orchestrated via an instruction-driven execution model. It integrates both a TPU for compute-intensive tasks (e.g., conv, see Fig. 4(b)) and a TMU for TM operators. The TPU features a systolic array comprising 128 threads, each consisting of 32 PEs. Each PE is equipped with a 9-bit signed multiplier. Input tensors are fetched from DRAM through an AXI4 interconnect, buffered, and broadcast to the threads. The commit buffer aggregates the results from all threads and transfers them back to DRAM. In contrast, the TMU is optimized for *data-movement-intensive* and *compute-light* operators. It reshapes, rearranges, or redirects datastreams retrieved from DRAM, and either writes them back to memory or forwards them to other computational threads. Forwarded datastreams can support element-wise operations (e.g., Add) or non-regular TM operators such as PixelShuffle (see Fig. 4(b)).

To maximize the TMU's efficiency in handling these operators, we employ several system-level strategies that enhance both performance and flexibility.

1) Tensor Prefetch and Output Forwarding: As shown in Fig. 5(b), a tensor prefetching strategy is employed to minimize off-chip memory access latency and enhance TM

efficiency. Two on-chip tensor buffers and two TMUs are configured in a double-buffering arrangement, where one buffer processes data while the other concurrently loads or stores datastreams to/from external DRAM. In TMU–TPU cooperative scenarios, a ping-pong mechanism enables the pre-scheduling of partially committed tensors from the TPU, effectively overlapping memory transfers with computation to mask TMU latency.

To further enhance pipeline concurrency, an output forwarding strategy is employed alongside prefetching. As shown in Fig. 5(c), when the TPU reaches the final stages of a compute-intensive operation (e.g., Conv, as shown in Fig.4(b)), it begins streaming partial output tensors to the buffer before completing the entire computation. This allows the TMU to initiate subsequent operators—such as PixelShuffle, or Add—early, thereby reducing idle cycles and improving throughput.

2) Block-based Manipulation: Coarse-grained TM operators operate on tensor blocks whose channel dimensions align with the burst width of the AXI interface. The TMU directly manipulates such datastreams and reshapes them within on-chip buffers, thereby increasing data throughput and minimizing memory latency. For fine-grained TM operators, a reconfigurable masking engine (RME) is introduced to optimize data paths and memory utilization, mitigating the performance penalties typically associated with sub-word data manipulation.

B. TMU's Microarchitecture and Dataflows

As illustrated in Fig. 6, the TMU integrates the generic execution model, address generator, reconfigurable masking engine (RME), and supporting control modules within a unified architecture.

A centralized finite-state machine (FSM) orchestrates the execution stages defined in Fig. 3, directing the instruction flow and coordinating data movement across the pipeline. To support the heterogeneous characteristics of TM operators, the FSM enables three configurable dataflows corresponding to fine-grained, coarse-grained, and element-wise processing.

- *Fetch and Decode:* The TMU retrieves instructions from local memory and decodes them to determine the operator type and operand configuration (see Fig. 6(c) ❶, ❷).

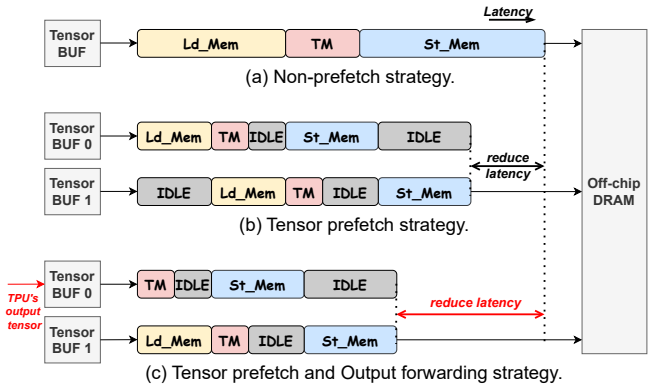

Fig. 5. (a) Non-prefetch strategy. (b) Tensor prefetch strategy. (c) Tensor prefetch and Output forwarding strategy.

- *Tensor Load:* Input data is fetched via a high-speed bus and stored in on-chip buffers. After potential reshaping or reordering, the datastream is moved to the commit buffer for further processing or storage (see Fig. 6(a)(b)(c) ❸).
- *Fine-grained TM:* This stage performs byte-level manipulation, such as selecting specific bytes (e.g., Rearrange, Bboxcal) or assembling data across segments (e.g., Transpose) to form new datastreams (see Fig. 6(a) ❹).
- *Coarse-grained TM:* This stage handles block-level reshaping operations driven by burst transfers (e.g., 16-byte AXI transactions). The address generator dynamically computes destination addresses to facilitate high-throughput tensor reorganization (see Fig. 6(b) ❻).
- *Element-wise Processing:* This stage supports arithmetic operations such as vectorized Add, Sub, and Mul. The TMU either executes these computations directly or cooperates with processing elements by preprocessing the datastream and forwarding it to global buffers (see Fig. 6(c) ❺).
- *Tensor Store:* The transformed datastream is written back from the commit buffer to off-chip DRAM or shared on-chip memory (see Fig. 6(a)(b)(c) ❼).
- *Branch:* For long tensors spanning multiple iterations, this stage updates memory addresses and buffer offsets to fetch subsequent segments, ensuring uninterrupted execution (see Fig. 6(c) ❽).

1) Address Generator: As illustrated in Fig. 7(a), the address generator implements the matrix-based addressing scheme defined by Eq. 1, enabling flexible computation of output memory addresses for coarse-grained TM operations. At runtime, the TMU instruction stream delivers operator-specific addressing parameters, which are decoded and loaded into dedicated configuration registers. These registers store the elements of matrices **A** and **B**, controlling the affine mapping from input coordinates (x_i, y_i, c_i) to output coordinates (x_o, y_o, c_o). The address generator executes the transformation in three pipeline stages. First, the row vectors of **A** are partitioned into three groups and multiplied with the input index vector. The resulting partial sums are then added to the corresponding entries in **B** to form an intermediate vector **C**. Finally, the output address $addr_{out}$ is computed by combining **C** with the base address $addr_{base}$ through post-addition logic. This output address is used to direct the storage of manipulated tensors to DRAM or to downstream modules. In conjunction with the write stride control, the generator iterates over tensor segments. Internal state registers are updated accordingly to determine whether the current TM instruction has been completed or requires additional iterations.

2) Reconfigurable Masking Engine: To realize fine-grained TM, the RME shown in Fig. 7(b) utilizes segment masking counters to acquire valid bus transfers into the tensor buffer. Afterwards, It supports two processing approaches for fine-grained TM, namely *assemble* and *evaluate*.

In the assemble scheme, identified bytes in the byte masking register are assembled into a new datastream in the assemble register, which is useful for operators like Rearrange, Rot90, and Transpose. The evaluate scheme processes selected bytes to extract result bytes, such as the case of Bboxcal, maximal,

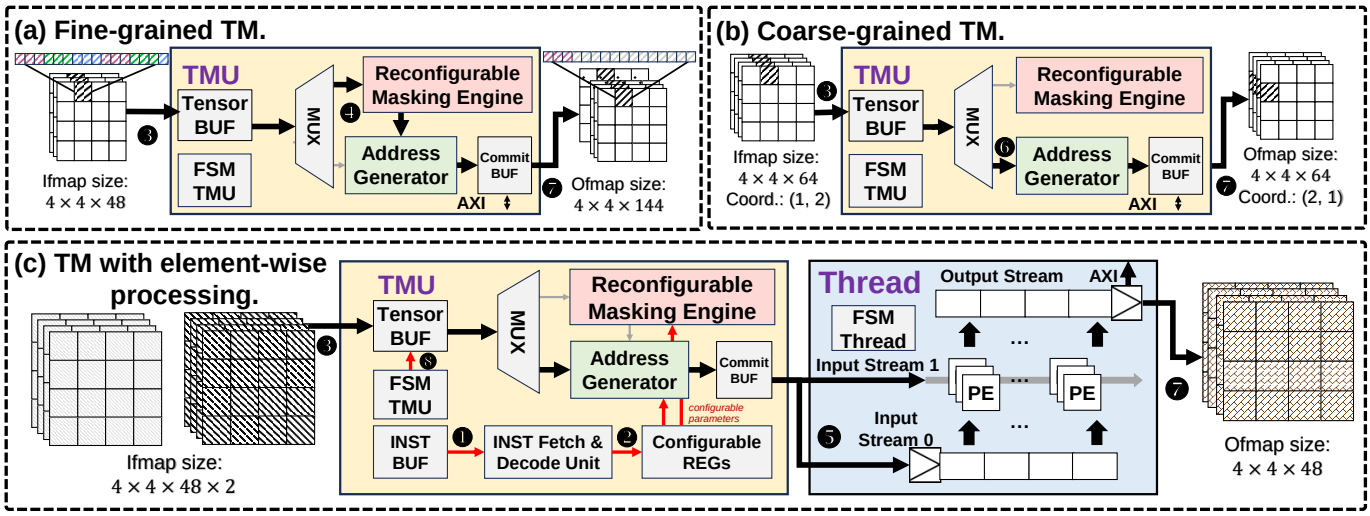

Fig. 6. TMU's microarchitecture and dataflows. (a) Fine-grained TM. (b) Coarse-grained TM. (c) TM with element-wise processing.

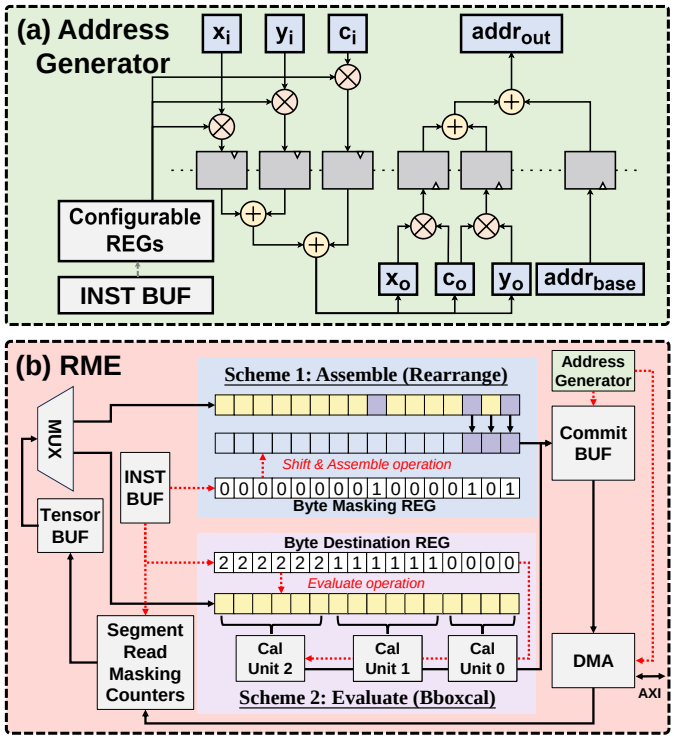

Fig. 7. (a) Address generator for coarse-grained TM operators. (b) Reconfigurable masking engine (RME) for fine-grained TM operators.

or minimal value retrieval from input datastream. The byte destination register maps bytes to specific calculation units, simplifying integration with computational logic.

While these fine-grained opeators are inherently operation-specific and difficult to build in a generalized structure, the RME mitigates this challenge by offering a structured abstraction for implementation. Specifically, both the *assemble* and *evaluate* schemes adhere to a predefined template comprising three stages: *(i) byte-level masking and indexing*, *(ii) construction and distribution of a new datastream*, and *(iii) conditional routing and commitment under FSM control*. For instance, a new TM op (such as selective value gating) can be mapped to this template by simply configuring its masking rule and output mapping logic, thereby avoiding the need for bespoke dataflow design.

In both schemes, the result datastream is output into the commit buffer under the direction of the address generator. Although both schemes may fill the commit buffer in an interleaved manner, after predictable rounds of processing, a renewed continuous datastream is formed in the commit buffer, which can be streamed to memories through DMA uninterruptedly. The timing of committing is controlled by the FSM tensor store stage.

TABLE III
OPERATOR CONFIGURATION PARAMETERS

TM Op.	Ifmap Size	Ofmap Size	Abbr.
Rearrange	448×448×3	448×448×16	RR
Resize	448×448×3	224×224×3	RS
Bboxcal	448×448×256	3×448×448×85	BC
Transpose	448×448×64	448×448×64	TS
Rot90	448×448×64	448×448×64	RT
Img2col	448×448×64	446×446×64	IC
PixelShuffle	448×448×64	896×896×16	PS
PixelUnshuffle	448×448×64	224×224×256	PU
Upsample	448×448×64	896×896×64	US
Route	2×448×448×64	448×448×128	RO
Split	448×448×64	2×448×448×32	SL
Add	448×448×64	448×448×64	AD

VI. Experimental Results

This section presents our evaluation methodology and summarizes the experimental results of the proposed accelerators in comparison with conventional CPU and GPU platforms.

A. Experimental Setup

1) Implementation: The proposed TMU and its coupled TPU are fully implemented in Verilog, verified through VCS simulation, and integrated into an in-house SoC infrastructure. The complete system is deployed on a Xilinx Kintex-7 410T

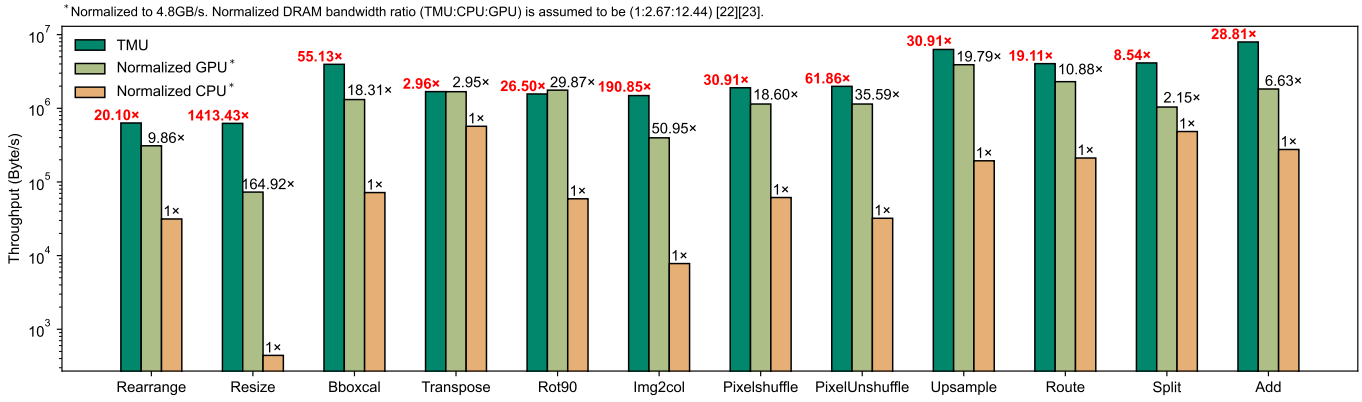

Fig. 8. Inference latency benchmarking (TMU Prefetch vs Normalized GPU vs Normalized CPU) for TM operators.

TABLE IV
MODEL CONFIGURATION PARAMETERS

DNN Type	DNN Model	Ifmap Size	TM Op.
CNN	ESPCN	448×448×3	RR, PS
	EDSR		RR, PS, AD
	YOLOv3		RR, RO, US, AD, BB
	YOLOv3-Tiny		RR, RO, US, BB
	YOLOv8	640×640×3	RR, RO, US, AD, SL, BB
Transformer	Attention	64×768	TS, RO

FPGA board, which communicates with the host environment via python-based interfaces and a USB 3.0 data link. Logic synthesis is performed using Synopsys Design Compiler with SMIC 40 nm low leakage standard cell libraries, targeting a clock frequency of 300 MHz. Fig. 9 shows the physical hardware prototype, featuring a camera-equipped FPGA board running real-time object detection with YOLOv8 [10], along with a PC-based software interface for system interaction and visualization.

2) Software Platform Setup: The evaluated neural networks are processed using an in-house AI compilation toolchain built with Python 3.6, TensorFlow 2.6, and NumPy 2.2. This toolchain parses each model to extract its computational graph and operator dependencies, applies post-training quantization (PTQ) using TensorFlow Lite to convert weights to INT8 precision, and transforms the result into an intermediate representation (IR) for hardware-specific optimization. The optimized models are then deployed to the target platforms. Using this toolchain, we conduct a comprehensive inference latency evaluation at both the TM operator and application levels across TPU, TMU, Jetson TX2, and Raspberry Pi 4.

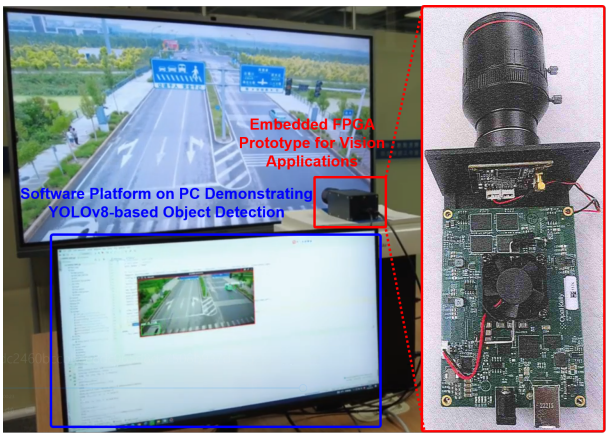

Fig. 9. Demonstration of a camera-equipped FPGA hardware prototype and PC-based software platform for YOLOv8 object detection

B. Performance Evaluation

1) TM operators benchmarking: We evaluated the performance of various TM operators on the TMU, comparing it to a 1.3GHz NVIDIA Pascal GPU (Jetson TX2) and a 1.5GHz ARM Cortex-A72 CPU (Raspberry Pi 4 Model B) using official TensorFlow library functions. The configuration parameters of TM operators are listed in Table III. The Jetson TX2 and Raspberry Pi 4 Model B were chosen for their relevance in edge computing, providing high programmability for executing TM operators, unlike other typical accelerators, which often lack support for a wide range of TM operators, limiting their applicability in such tasks. Additionally, both devices employ a complex multi-level caching mechanism for data movement, which contrasts with the near-memory DMI approach adopted by the TMU.

Note that the performance of TM operators is heavily constrained by DRAM bandwidth. To enable fair comparison across platforms, the measured performance of the CPU (with a DRAM bandwidth of 12.8GB/s [13]) and the GPU (with a bandwidth of 59.7GB/s [12]) is normalized to match the DRAM bandwidth of the TMU, which is 4.8GB/s. This normalization ensures that observed performance differences reflect architectural design efficiency rather than bandwidth disparities, enabling a more meaningful and bandwidth-fair comparison across platforms.

As shown in Fig. 8, the TMU demonstrates substantial throughput improvements, It achieves up to 1413.43× in Resize and 61.86× in PixelUnshuffle than CPU, consistently outperforming the normalized GPU and CPU. The TMU excels in various fine-grained TM operators, such as 55.13× in Bboxcal. Significant gains are also observed in element-wise operators like 19.11× in Route and 28.81× in Add.

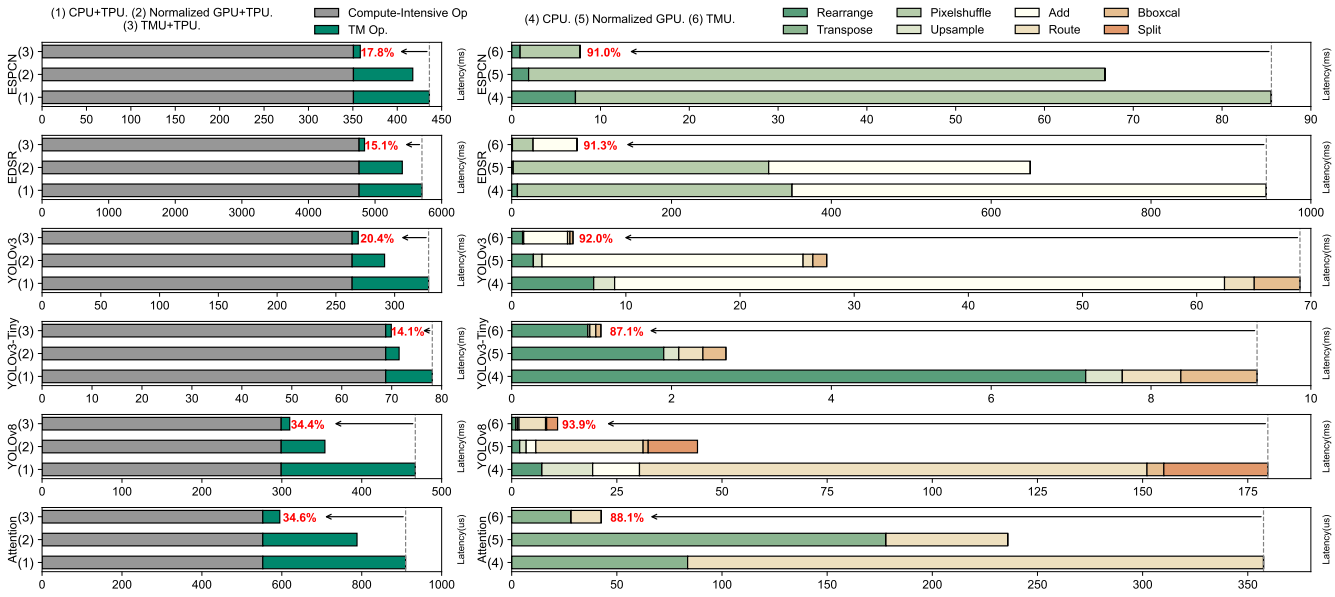

(a) Inference latency benchmark for the entire neural network.

(b) Inference latency benchmark for TM operators only.

Fig. 10. Graphical representation of typical TM operators adopted in state-of-the-art neural networks.

These results highlight the TMU's potential in performance-boosting of TM operators. The only TM operator where the TMU underperforms GPU is Rot90, due to the time-intensive data disassembling and reassembling, particularly between the width and channel dimensions, which can be further optimized in the ongoing TMU implementation.

2) Application-level benchmarking: To ensure fairness and reproducibility in application-level benchmarking, we adopt a unified evaluation methodology. On CPU and GPU platforms, TensorFlow implementations are instrumented with custom profiler context blocks to measure the execution time of TM layers. All reported results exclude cross-platform communication overhead to isolate computation performance. In this setup, compute-intensive operators, including convolutions and associated transformations (e.g., *Img2col*), are executed on the TPU, while TM operators are handled by the TMU, CPU, and GPU platforms. To highlight the TMU's contribution to neural network acceleration, the impact of *Img2col* is excluded from the application-level benchmarking.

Execution times for six typical deep learning networks featuring various TM operators are shown in Fig. 10(a). These applications include ESPCN [25], EDSR [26], YOLOv3 [27], YOLOv3-Tiny [28], YOLOv8 [10], and Attention [9]. The configuration parameters of application-level benchmarking are listed in Table IV. Combined with our in-house TPU [11], the TMU outperformed a conventional CPU (even without performance normalization) in inference latency benchmark for the entire neural network, achieving speedup of 17.8%, 15.1%, 20.4%, 14.1%, 34.4%, and 34.6% for ESPCN, EDSR, YOLOv3, YOLOv3-Tiny, YOLOv8, and Attention, respectively. Fig. 10(b) shows the accumulated latency for all TM operators, where the TMU demonstrated significant reductions in latency, achieving reductions of 91.0%, 91.3%, 92.0%, 87.1%, 93.9%, and 88.1% for ESPCN, EDSR, YOLOv3, YOLOv3-Tiny, YOLOv8, and Attention, respectively.

C. Physical Overheads

Table V provides a comparison between the proposed TMU and several state-of-the-art DMI accelerators. It is important to note that the reported physical parameters do not include the DRAM controller or interface circuitry. Compared to prior works, this design is the first DMI accelerator to serve as a generic design template with reconfigurable registers capable of supporting a wide range of TM operators. In addition, the proposed TMU demonstrates clear advantages in both area and power efficiency. After normalization to a 40 nm process, it achieves the smallest area footprint (0.019 mm^2), which is 15.3× smaller than that of AME [29]. While the power comparison is based on reported values without frequency normalization, the TMU still consumes 1.52× less power (2.7 mW vs 4.1 mW). Note that the adopted normalized area and ratio refers to paper [14] and report [15].

Furthermore, this work supports a broader range of functions (①-⑧), while AME [29] and ECNN [30] lacks functional support. These advantages underscore the efficiency of the proposed architecture in minimizing power consumption and area footprint while maintaining competitive performance and extensive functional support. Additionally, The TMU+TPU combination achieves high-performance improvement and a series of TM operators with only 0.07% extra overhead of the TPU's area, this indicates that the TMU contributes significantly to system-level optimization.

VII. Conclusion

This work presents TMU, a reconfigurable near-memory TMU designed to accelerate DMI operators in modern AI

TABLE V
COMPARISON WITH STATE-OF-THE-ART DATA-MOVE-INTENSIVE ACCELERATORS.

Accelerators	ECNN [30]†	AME [29]	This Work
Technology	40 nm	7 nm	40 nm
Frequency(MHz)	250	2100	300
Area(mm^2)	2.26	0.034	0.019
Power(mW)	100	4.1	2.7
Normal. Area‡(mm^2)	2.26	0.291	0.019
Reconfigurability	✘	✘	✔
Function§	④,⑥,⑨	N/A	①-⑧

	Freq.	Area	Power	DRAM	#.MACs
Integrated TPU [11]	300 MHz	26.96 mm^2 post P&R	1.83 W w. DRAM	1200 MT/s DDR3	4096 (int8)

†ECNN [30] includes Src/SrcS/Dst Reorder, ADDE/ACCI, MUXA/MUXC, Block Buffer File, and other modules.
‡Normalized area ratio (7 nm:40 nm) = (1:8.57) [14], [15]. Normalized to 40 nm, 300 MHz.
§①: Rearrange, ②: Resize, ③: Bboxcal, ④: Rot90 and Transpose, ⑤: Img2col, ⑥: PixelShuffle and PixelUnshuffle, ⑦: Upsample, ⑧: Route, Split, and Add, ⑨: Downsample.

workloads. By addressing the frequently neglected latency bottlenecks introduced by TM operators, TMU significantly enhances system throughput across multiple state-of-the-art models. Despite its compact hardware footprint, the TMU delivers substantial performance gains and demonstrates strong scalability and integration potential within high-throughput AI SoCs.

REFERENCES

[1] A. Ivanov, N. Dryden, T. Ben-Nun, S. Li, and T. Hoefler, "Data movement is all you need: A case study on optimizing transformers," *Proceedings of Machine Learning and Systems*, vol. 3, pp. 711–732, 2021.

[2] M. Khairy, A. G. Wassal, and M. Zahran, "A survey of architectural approaches for improving gpgpu performance, programmability and heterogeneity," *Journal of Parallel and Distributed Computing*, vol. 127, pp. 65–88, 2019.

[3] F. Tu, Y. Wang, Z. Wu, W. Wu, L. Liu, Y. Hu, S. Wei, and S. Yin, "16.4 tensorcim: A 28nm 3.7 nj/gather and 8.3 tflops/w fp32 digital-cim tensor processor for mcm-cim-based beyond-nn acceleration," in *2023 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, 2023, pp. 254–256.

[4] W.-Y. Hsu and W.-Y. Lin, "Ratio-and-scale-aware yolo for pedestrian detection," *IEEE transactions on image processing*, vol. 30, pp. 934–947, 2020.

[5] N. Brown and D. Dolman, "It's all about data movement: Optimising fpga data access to boost performance," in *2019 IEEE/ACM International Workshop on Heterogeneous High-performance Reconfigurable Computing (H2RC)*, 2019, pp. 1–10.

[6] M. Tang and S. Liu, "A dynamic computational memory address architecture for systolic array cnn accelerators," in *2022 IEEE 24th Int Conf on High Performance Computing & Communications; 8th Int Conf on Data Science & Systems; 20th Int Conf on Smart City; 8th Int Conf on Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys)*, 2022, pp. 1314–1319.

[7] A. Rawal, Y. Fang, and A. Chien, "Programmable acceleration for sparse matrices in a data-movement limited world," in *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2019, pp. 47–56.

[8] H. Liao, J. Tu, J. Xia, H. Liu, X. Zhou, H. Yuan, and Y. Hu, "Ascend: a scalable and unified architecture for ubiquitous deep neural network computing: Industry track paper," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2021, pp. 789–801.

[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[10] R. Varghese and M. Sambath, "Yolov8: A novel object detection algorithm with enhanced performance and robustness," in *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*. IEEE, 2024, pp. 1–6.

[11] Y. Li, Z. Wang, W. Ou, C. Liang, W. Zhou, Y. Yang, and C. Chen, "Low-latency buffering for mixed-precision neural network accelerator with multap and fqpipe," in *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2024, pp. 1–5.

[12] NVIDIA, "Jetson tx2 series module datasheet v1.8," 2022. [Online]. Available: https://openzeka.com/wp-content/uploads/2022/07/Jetson-TX2-Series-Module-Datasheet-v1.8.pdf

[13] K. M. Huynh, T. T. Bui Nguyen, H. V. Nguyen, K. Dac Tran, K. Iwata, K. Mizumoto, N. Honda, K. Matsumoto, K. Matsubara, and S. Mochizuki, "16.8 gb/s lpddr4-3200@32-bit memory access bandwidth," in *2017 7th International Conference on Integrated Circuits, Design, and Verification (ICDV)*, 2017, pp. 16–21.

[14] H. Mo, W. Zhu, W. Hu, Q. Li, A. Li, S. Yin, S. Wei, and L. Liu, "A 12.1 tops/w quantized network acceleration processor with effective-weight-based convolution and error-compensation-based prediction," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 5, pp. 1542–1557, 2021.

[15] TSMC. (2019) Tsmc 2019 annual report. [Online]. Available: https://investor.tsmc.com/static/annualReports/2019/english/ebook/index.html

[16] S. Gomar, M. Mirhassani, and M. Ahmadi, "Precise digital implementations of hyperbolic tanh and sigmoid function," in *2016 50th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2016, pp. 1586–1589.

[17] Z. Pan, Z. Gu, X. Jiang, G. Zhu, and D. Ma, "A modular approximation methodology for efficient fixed-point hardware implementation of the sigmoid function," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 10, pp. 10 694–10 703, 2022.

[18] N. P. Jouppi, C. Young, N. Patil, D. Patterson, and G. A. E. Al, "In-datacenter performance analysis of a tensor processing unit," in *Computer architecture news*, 2017, pp. 1–12.

[19] Y. Chen, Y. Xie, L. Song, F. Chen, and T. Tang, "A survey of accelerator architectures for deep neural networks," *Engineering*, vol. 6, no. 3, pp. 264–274, 2020.

[20] T. E. Oliphant *et al.*, *Guide to numpy*. Trelgol Publishing USA, 2006, vol. 1.

[21] P. Diniz and J. Park, "Data reorganization engines for the next generation of fpgas," in *Proc. of the ACM Conf. on Field-Programmable-Gate-Arrays (FPGA'02)*, 2002, pp. 100–110.

[22] K. Zhang, X. Zhang, and Z. Zhang, "Tucker tensor decomposition on fpga," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–8.

[23] S. Lloyd and M. Gokhale, "In-memory data rearrangement for irregular, data-intensive computing," *Computer*, vol. 48, no. 8, pp. 18–25, 2015.

[24] B. Akin, F. Franchetti, and J. C. Hoe, "Data reorganization in memory using 3d-stacked dram," *ACM SIGARCH Computer Architecture News*, vol. 43, no. 3S, pp. 131–143, 2015.

[25] M. A. Talab, S. Awang, and S. A.-d. M. Najim, "Super-low resolution face recognition using integrated efficient sub-pixel convolutional neural network (espcn) and convolutional neural network (cnn)," in *2019 IEEE international conference on automatic control and intelligent systems (I2CACIS)*. IEEE, 2019, pp. 331–335.

[26] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, "Enhanced deep residual networks for single image super-resolution," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 136–144.

[27] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[28] P. Adarsh, P. Rathi, and M. Kumar, "Yolo v3-tiny: Object detection and recognition using one stage improved model," in *2020 6th international conference on advanced computing and communication systems (ICACCS)*. IEEE, 2020, pp. 687–694.

[29] Y. Sugawara, D. Chen, R. A. Haring, A. Kayi, E. Ratzlaff, R. M. Senger, K. Sugavanam, R. Bellofatto, B. J. Nathanson, and C. Stunkel, "Data movement accelerator engines on a prototype power10 processor," *IEEE Micro*, vol. 43, no. 1, pp. 67–75, 2022.

[30] C.-T. Huang, Y.-C. Ding, H.-C. Wang, C.-W. Weng, K.-P. Lin, L.-W. Wang, and L.-D. Chen, "ecnn: A block-based and highly-parallel cnn accelerator for edge inference," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 182–195.