

MACO: Exploring GEMM Acceleration on a Loosely-Coupled Multi-core Processor

Bingcai Sui^{1,2}, Junzhong Shen^{1,2}, Caixia Sun^{1,2}, Junhui Wang^{1,2}, Zhong Zheng^{1,2}, Wei Guo^{1,2}✉

¹*School of Computer, National University of Defense Technology*

²*Key Laboratory of Advanced Microprocessor Chips and Systems
Changsha, China*

{bingcaisui, shenjunchong, cxsun, wangjunhui, zheng-zhong, wineer-guowei}@nudt.edu.cn

Abstract—General-purpose processor vendors have integrated customized accelerator in their products due to the widespread use of General Matrix-Matrix Multiplication (GEMM) kernels. However, it remains a challenge to further improve the flexibility and scalability of these GEMM-enhanced processors to cater to the emerging large-scale GEMM workloads. In this paper we propose *MACO*, a novel loosely-coupled multi-core general-purpose architecture optimized for GEMM-related applications. To enhance the programmability and flexibility of *MACO*, the paper introduces a tile-based instruction set architecture. Additionally, the paper presents techniques such as hardware-assisted data prefetching and locking, and predictive address translation to further enhance the computational efficiency of *MACO* for GEMM workloads. The experimental results demonstrate that *MACO* exhibits good scalability, achieving an average computational efficiency of 90% across multiple cores. Furthermore, evaluations on state-of-the-art deep neural networks show that *MACO* can achieve up to 1.1 TFLOPS with 88% computational efficiency, indicating its adaptivity to deep learning workloads.

Index Terms—general-purpose processor, GEMM, loosely-coupled architecture, tile-base instruction set

I. INTRODUCTION

General Matrix-Matrix Multiplication (GEMM) is a critical building block for many domains including, but not limited to, high-performance computing (HPC), computer vision (CV), and natural language process (NLP). The demand for high-efficiency GEMM computations has driven the development of domain-specific architectures (DSA) such as Google’s TPU [1] and Nvidia’s tensor core, which are designed to meet the performance and energy-efficiency requirements. To address the need for efficient AI application processing, CPU vendors like Intel, IBM, and ARM have started integrating deep learning-specific co-processors or execution units in their products. These solutions, referred to as GEMM-enhanced CPUs in the paper, extend the Instruction Set Architectures (ISAs) of Intel (AMX) and ARM (SME) to enable CPUs to execute GEMM workloads on customized execution units.

The paper classifies GEMM-enhanced CPUs into two categories: tightly-coupled architectures (TCA) [2], [3] and loosely-coupled architectures (LCA) [4], [5]. In TCA, the matrix accelerators are considered as an integral part of the CPU’s pipeline. On the other hand, LCA designs treat the matrix accelerators as co-processors for the CPU core. The authors highlight that while TCA offers advantages like reduced area budget and synchronization overhead between the CPU core and matrix accelerator, the performance of TCA solutions can

be impacted due to resource contention between the CPU core and matrix accelerator.

LCA solutions provide several advantages compared to tightly-coupled architectures. One of the key benefits is the ease of use they offer. Loosely-coupled architectures simplify the design and implementation process, making it more straightforward to integrate and utilize components like the CPU and matrix accelerator. Furthermore, these architectures excel in parallel computing, enabling efficient parallelization of tasks between the CPU and matrix accelerator. In this way, LCA solutions have a wider range of application scenarios, for example, when deploying recommended system on these architectures, we can offload top and bottom MLPs to the matrix engine leaving the CPU core free to run embedding lookups.

Overall, loosely-coupled architectures contribute to enhanced usability and facilitate parallel computing between the CPU and matrix accelerator, leading to improved performance in various applications. However, the generality of the loosely-coupled architectures can be further improved, and this is the main focus of our paper. Gemini [4], a representative of loosely-coupled architectures, provides address translation support but does not consider the possible overhead of the accelerator in memory access caused by frequent cache misses when dealing with large-scale GEMM workloads. Additionally, Gemini does not provide a specific solution for multi-process support and exception handling. Telum [5] supports multi-processing and exception event handling, but the processes on multiple cores cannot run GEMM in parallel since only one process can occupy the shared AI accelerator at a time (via arbitration), causing frequent process switching and recovery. Moreover, although Gemini and Telum both implement multi-core solutions, none of them provide details of mapping GEMM tasks on multiple cores, and the benefits of parallel computing on multiple cores are not being fully exploited.

To address the above issues, this work proposes *MACO*, a loosely coupled multi-core general-purpose processor architecture with enhanced GEMM computation ability. Compared with Gemini, *MACO* provides better hardware support in multi-process execution, virtual-to-physical address translation, and exception event handling. Compared to Telum, all general-purpose cores of *MACO* are accompanied by a separate matrix multiplication acceleration engine (*MMAE*), allowing all cores

to execute GEMM processes simultaneously, providing a higher degree of parallelism for GEMM. In addition, to improve *MACO*'s usability and adaptability to various GEMM-related applications, this work optimizes the architectural design from the following aspects: (1) An instruction set is proposed to provide users with a variety of functions including GEMM computation, data migration, and data prefetch; (2) Efforts are made to enhance the architectural support for address translation, multi-process management and exception handling. Moreover, We further explore the possibility of parallel computation of CPUs and MMAEs on *MACO* for applications combined with both GEMM and non-GEMM workloads. Finally, this work goes deeper than state-of-the-art Matrix-Multiplication-enhanced CPUs by presenting details of mapping scheme for GEMM workloads on multiple cores of *MACO*. The main contributions of this work are summarized as follows:

- We propose a novel loosely-coupled multi-core processor architecture named *MACO*, which is featured by integrating multiple CPU+MMAEs (GEMM Acceleration Engine) as well as a highly scalable NOCs with cache coherence.
- An extended instruction set is proposed to improve the programmability and flexibility of *MACO*, which also exposes *MACO*'s rich functionality to users.
- We developed latency-hiding address translation technique based on page table address prediction to further improve GEMM performance on *MACO*.

The experimental results indicate that *MACO* has the capability to achieve a maximum throughput of 1.1 TFLOPS while maintaining a high computational efficiency of 88%.

II. BACKGROUND

A. GEMM-enhanced CPUs

Tightly-coupled architectures. Like Intel's Advanced Matrix Extensions (AMX) [2], RASA [3] which place MAE inside the CPU core, the main feature of tightly-coupled architecture (TCA) is regarding MAE as part of the CPU pipeline. The obvious advantage of TCA is that the MAE can share the resources of the CPU core, contributing to a reduction in the area overhead of the CPU chip. However, the execution procedures of the CPU and MAE are also tightly-coupled, meaning that both of them may suffer performance loss due to competition for resources (e.g. MMU, LSU).

Loosely-coupled architectures. Representative examples for loosely-coupled architectures (LCA) include Telum and Gemini, which consider MAE as a co-processor of the CPU core. In this scenario, since MAE has independent data paths and data load/store units (e.g. DMA engine) for accessing last level cache or external memory, providing better opportunities for parallel computing between CPU and MAE. However, the main defects of LCA lie in the following aspects: (1) high synchronisation overhead between the CPU core and MAE; (2) MAE has difficulties in performing address translation independently; (3) process management and exception handling become challenges for the CPU core.




Fig. 1. Illustration of mapping tile GEMM algorithm on systolic array.

B. Mapping Tile GEMM algorithm on Systolic Arrays

Systolic arrays [6] have been successfully implemented in various commercial products, including Intel's AMX [2], Google's TPU [1], and IBM's Telum [5]. This is due to their advantages in simple construction, high concurrency, and efficient exploitation of the inherent data reuse of algorithms. Typical systolic arrays consist of many homogeneous processing elements (PE), each responsible for a MAC operation and interacting with each other through localized short concatenation lines for data interaction.

The left part of Fig. 1 shows the classical tiled GEMM algorithm which breaks up the large matrices A, B and C into smaller sub-matrices for efficient computation. The right part of Fig. 1 illustrates the schematic diagram of the equivalent computation procedure of tiled-GEMM algorithm mapped on a systolic array that employs the input-stationary data flow. In this case, the data of the sub-matrix B is pre-loaded and buffered within the systolic array before the data of sub-matrices A and C stream in. As each PE receives its data from the sub-matrices, it performs a local MAC operation and forwards the partial products either vertically (in the column direction) to its neighboring PE. During the computation, the partial products are temporarily stored in on-chip buffers, allowing for efficient data handling. These partial products are then loaded back into the systolic array to facilitate subsequent computations. This repeated procedure continues until the final results of the matrix multiplication operation are obtained.

III. ARCHITECTURAL DESIGN OF *MACO*

A. Overview

In Fig. 2, *MACO* consists of up to 16 homogeneous compute nodes interconnected by a network on chip (NOC). Each compute node is integrated with a general-purpose processor core that is associated with a matrix multiplication acceleration engine (*MMAE*). The *MMAE* shares the CPU core's shared TLB (sTLB) via customized interfaces. However, the CPU's L1 data cache and L2 cache are not accessible to *MMAE*. Similarly, the CPU has no access to the associated *MMAE*'s on-chip buffers. The CPU core is a 64-bit high-performance general-purpose processor core with a multi-issue superscalar architecture. Table I reports the architectural parameters of a CPU core.

Form the right part of Fig. 2(a), it can be seen that *MMAE* is built on a 4x4 two-dimensional systolic array (SA) with integrated high-capacity buffers of 192KB for efficient data reuse.

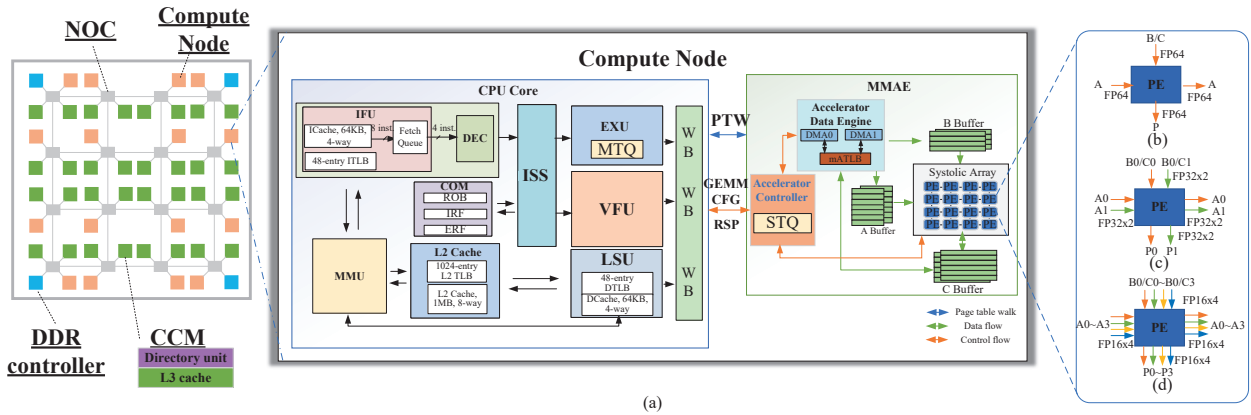


TABLE I
ARCHITECTURAL PARAMETERS OF A CPU CORE

Its main function is to execute tile GEMM operations. More specifically, the integrated Accelerator Data Engine (ADE) is responsible for transferring data between L3 cache and on-chip buffers. And the Accelerator Controller (AC) functions by receiving configurations from the CPU and then scheduling SA, ADE and AC modules to complete the GEMM tasks. Different from previous work, We further extend the classical dataflow of systolic array (see Fig. 1) to support SIMD-like compute modes including 2-way FP32 (Fig. 2(c)) and 4-way FP16 (Fig. 2(d)) parallel computations. In addition, by integrating powerful DMA engines, *MMAE* can carry out high-capacity data initialization and data migration without disturbing the CPU core. More importantly, the above procedures are fully programmable via the proposed instructions (details follow), which effectively improves the flexibility of *MMAE*.

B. Matrix Processing Assist Instruction Set

C. Multi-process Management

TABLE II
ILLUSTRATIONS OF THE PROPOSED INSTRUCTION SET.

Functions	Instructions	Description	Usage
Data migration	MA_MOVE	Copy data from source address to destination address.	MA_MOVE Rd, Rn
	MA_INIT	Set data in destination space to zeros.	MA_INIT Rd, Rn
	MA_STASH	Perform data prefetch from the external memory to L3 cache.	MA_STASH Rd, Rn
GEMM computing	MA_CFG	Request an <i>MTQ</i> entry for executing a GEMM task.	MA_CFG Rd, Rn
Task management	MA_READ	Obtain the execution state of a certain GEMM task.	MA_READ Rd, Rn
	MA_STATE	Obtain execution state of a certain GEMM task and release the occupied <i>MTQ</i> entry.	MA_STATE Rd, Rn
	MA_CLEAR	Clear a certain <i>MTQ</i> entry.	MA_CLEAR, Rn

TABLE III
DETAILS OF AN *MTQ* ENTRY.

Field	Description
Valid	Indicate whether the entry is allocated.
Done	Indicate whether the task is completed.
ASID	Process identifier.
exception_en	Indicate exception occurs during <i>MMAE</i> 's task execution.
exception_type	Specific type of an exception event.




Fig. 3. State transition diagram of an *MTQ* entry.

of an *MTQ* entry. It can be seen that an *MTQ* entry can provide rich information including the process identifier, task execution state and exception events. The functions of *STQ* include: receiving parameters of a GEMM task from the CPU core (identified by the same MAID), parsing parameters and saving them at its local registers, monitoring execution states of other components of *MMAE* (e.g. DMA engines, systolic arrays), and responding the status of the GEMM task to the corresponding *MTQ* entry.

For GEMM workloads, the CPU core would send the GEMM relative parameters of the task to the *MMAE*, thereby the *STQ* entry specified by the MAID would receive and then buffer the configuration information locally. The buffered tasks in the *STQ* entries will be automatically executed when the active *STQ* entry has completed its task. Fig. 3 shows the state transition diagram of an *MTQ* entry when process switch occur. It can be seen that we can combine the values of "Done" and "ASID" from the specified *MTQ* entry to determine the execution state of the process, even though the entry has been occupied by other process (see state 3). For state 4, since a GEMM task may be automatically terminated by the *MMAE* if there are exception events during task execution, users have to do further check to determine the accurate type of the exception events. Note that both *MTQ* and *STQ* will not be affected by process switching, thereby we can obtain reliable information associated with all the processes from their corresponding allocated *MTQ* entries.

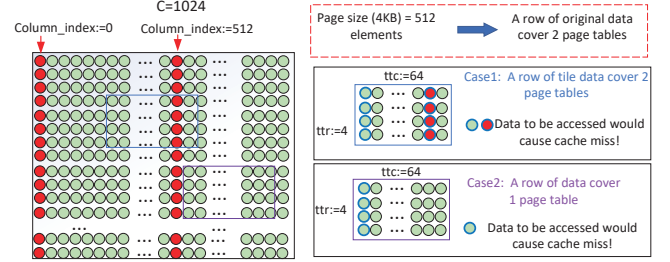


Fig. 4. Basics of page table address prediction.

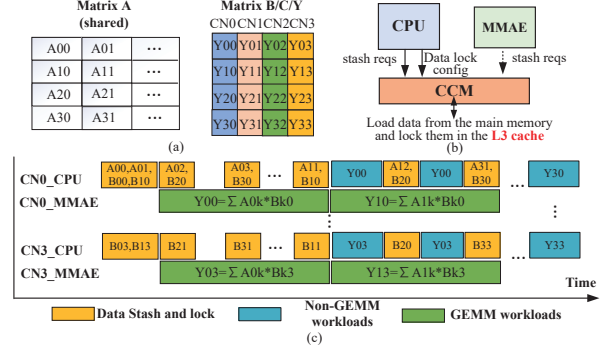


Fig. 5. Mapping schemes of GEMM⁺ workloads on MACO.

IV. IMPLEMENTATION DETAILS OF *MACO*

A. Predictive Address Translation

Fig. 4 presents a simple example to illustrate how to determine the virtual addresses that would cause cache misses. Note that parameters such as the number of columns of the original matrix (*C*), the tile size (i.e. $\langle Tr, Tc \rangle$), and the page table size are configured to the *MMAE* in advance. In Fig. 4, we assume the page table size is 4KB, and the original matrix consists of 1024 elements in FP64 precision (8KB in total), meaning that each row of the original matrix data are mapped to two page tables. As it can be seen that the elements identified by red circles represent the first data located at each page table, and once the location of the tile data in the original matrix is known, we can determine whether the data to be accessed would cause cache miss. Based on this observation, we design a module named *mATLB* to generate multiple virtual addresses in advance, then sends them to the CPU core's memory manage unit (MMU) to perform page table walk. After a period of time, the returned *ATLB* entries (including translated physical address) would be stored in the local buffers of *mATLB*. Each entry would be accessed by the DMA engines then provide physical address for their memory access requests, and each entry would be removed from the buffer once it fails to match the current virtual address. In this way, the overhead of cache misses would be perfectly hidden by pre-performed page table walk.

B. Mapping Real-world GEMM⁺ Workloads on MACO

In real-world applications involving CNNs, RNNs and transformer-based models, it is common to follow GEMM-based workloads (such as convolutional layers, fully-connected layers, and attention layers) with non-GEMM but complex workloads, including normalization, activation, and softmax functions. This work present a novel and efficient scheme

TABLE IV
COMPARISONS OF THE CPU CORE AND *MMAE*.

	Freq (GHz)	Area (mm ²)	Power (W)	FMACs	Peak Perf ^a (GFLOPS)
CPU	2.2	6.25	2.0	8	35.2(FP64)/71(FP32)
<i>MMAE</i>	2.5	1.58 ^b	1.5	16	80(FP64)/160(FP32)/ 320(FP16)

^a Theoretical peak performance, calculated by $2 \times \text{Freq} \times \text{FMACs}$.

^b Area breakdown: Buffers: 36.7%, SA: 24.7%, AC: 23.4%, ADE: 15.8%.

for mapping these combined GEMM (GEMM⁺ for short) workloads onto *MACO*. Fig. 5(b) provides an illustration of the data prefetch and locking procedure. Both the CPU and *MMAE* can issue "stash" requests to the CCM, enabling the prefetching of sub-matrices into the L3 Cache. Subsequently, the CPU can generate configurations lock the data in the L3 cache via the CCM. The advantages of data prefetching and locking two-fold: first, the memory access efficiency of *MMAE* could be guaranteed since there are no page misses occur during Page Table Walk (PTW); secondly, if the result tiles of the GEMM workloads also reside in the L3 cache, the CPU can perform subsequent non-GEMM workloads without incurring data misses. Fig. 5(c) illustrates the timing graph depicting the mapping of GEMM⁺ workloads onto the four compute nodes of *MACO*. As shown that we achieve achieve parallelization by tiling the original matrices and efficiently allocating the resulting sub-matrices to the compute nodes (CNs). (see Fig. 5(a)).

V. EVALUATIONS

A. Experimental Settings

To evaluate the *MACO* architecture, we performed both ASIC and FPGA flows. In the ASIC flow, we ensured that the timing requirements of the CPU cores, *MMAEs*, and NOC were met at frequencies of 2.2 GHz, 2.5 GHz, and 2.0 GHz, respectively, after Placement and Routing. These timings were achieved using a 12 nm library. In the FPGA flow, the RTL code of *MACO* underwent FPGA synthesis and layout using Vivado Design Suite 2022.2. The experiments were conducted on the Xilinx VCU440 FPGA, and the implemented *MACO* design was clocked at 50 MHz. On the FPGA platform, we successfully ran a modified Linux operating system on *MACO*. Note that GEMM workloads of various sizes used for evaluation were obtained from an open-source software package [7].

B. Experimental Results

1) *Evaluations on Area and Power*: Table IV compares the area and power consumption of a single CPU core and *MMAE*. It can be seen that the area of *MMAE* is only 25% of the size of CPU core, but the peak performance in GFLOPS of *MMAE* is over 2× of that of CPU core. As a result, *MMAE* can obtain a much higher (9×) area efficiency (GFLOPS/mm²) than CPU. In addition, The power consumption of *MMAE* is 25% lower than CPU, contributing to a 2× theoretical computation efficiency (GFLOPS/W) than CPU. It can be concluded that we have effectively extended the GEMM computation power of the CPU at a smaller cost of area and power consumption.




Fig. 6. Performance of *MACO* with/without page table prediction.




Fig. 7. Scalability of *MACO*.

2) *Evaluations on Address Translation*: To evaluate the effectiveness of the proposed address translation technique, we conducted experiments to measure the sustainable performance of *MACO* with and without prediction on address translation. In these experiments, we maintained a uniform page size and tiling size for all cases: 4KB pages and first-level tiling of $\langle Tr, Tc \rangle = \langle 1024, 1024 \rangle$, along with second-level tiling of $\langle ttr, ttc \rangle = \langle 64, 64 \rangle$. For simplicity, only one compute node was involved in these tests.

As depicted in Fig. 6, it is evident that *MACO* achieves higher computational efficiency when performing address translation in advance. Note that computational efficiency is calculated by the ratio of measured performance to theoretical peak performance in GFLOPS. The performance gap reaches a maximum of 6.5% with a matrix size of 1024. However, for matrix sizes smaller than 512, the performance gains are not significant (which is expected), amounting to less than 2%. This can be attributed to the fact that in those cases, the matrices fail to cover multiple page tables, resulting in fewer cache misses during the memory access of the *MMAE*. Consequently, the predictive address translation yields limited performance improvement in such scenarios.

3) *Scalability*: To assess the scalability of *MACO*, we conducted throughput tests using various GEMM workloads and varying the number of compute nodes (2, 4, 8, 16) involved. Each compute node was assigned an independent GEMM workload, with no inter-node interaction. The results, as depicted in Fig. 7, indicate the average computational efficiency per compute node (y-axis) across different matrix sizes (x-axis). It is observed that with an increase in the number of compute nodes, *MACO* experiences an average performance loss of 10%. This can be attributed to the NOC being unable to meet the bandwidth requirements of all compute nodes working in parallel. Despite the performance loss, *MACO* still achieves an




Fig. 8. Comparisons with state-of-the-art solutions with DL workloads.

approximate computation efficiency of 90% for all test cases. This result demonstrates that the proposed data prefetch and predictive address translation techniques effectively reduce the memory access overhead. Hence, we can conclude that *MACO* exhibits good scalability in parallel computing with multiple compute nodes.

4) *Comparisons with state-of-the-art*: To validate the superiority of *MACO*, we conducted experiments using real-world deep learning (DL) workloads. The benchmarks selected for this study were Resnet-50 [8], BERT [9], and GPT3 [10], all of which are used for inference purposes and employ FP32 precision. We compared *MACO* against four counterpart solutions, including (1) Baseline-1, *MACO* with CPU-only; (2) Baseline-2 *MACO* with MMAE, but without applying mapping scheme illustrated in Section IV.B; (3) A modified version of MacSim with similar configurations to RASA [3]; (4) Gemmini. To ensure a fair comparison, we configured all solutions with the same number of processing elements (16×16). As illustrated in Fig. 8, *MACO* outperformed the other solutions across all benchmarks. More specifically, *MACO* achieves an average performance gain of 1.35x and 1.30x over RASA and Gemmini, respectively. When compared to Baseline-1 and Baseline-3, *MACO* demonstrated remarkable performance improvements of 3.30x and 1.45x, respectively. In addition, *MACO* can achieve up to 1.1 TFLOPS with 88% computational efficiency, indicating its adaptivity to deep learning workloads.

VI. RELATED WORK

In the area of CPUs with efficient support for running dense or sparse GEMM operations, there are several related works that have been proposed.

For dense GEMM support on CPUs, ZCOMP [11] introduces a vector Instruction Set Architecture (ISA) extension that reduces cross-layer communication in Deep Neural Networks (DNNs). This extension aims to improve performance by optimizing the data movement between different layers of the DNN. RASA [3] proposes control and data optimizations for CPU matrix engines to improve utilization through efficient pipelining and overlap. It divides matrix multiplication into different sub-stages on the systolic array and introduces optimizations with pipelining and overlapping to maximize throughput.

In terms of sparse GEMM support on CPUs, much of the related work focuses on accelerating Deep Neural Networks (DNNs). SAVE [12] is a sparsity-aware CPU vector engine that skips redundant computations in sparse GEMM operations, thereby accelerating sparse DNN and high-performance computing workloads. SparCE [13] aims to increase the utilization

of vector engines by tracking general-purpose registers with zeros. It skips ineffective code blocks based on sparse input by annotating skippable code blocks in software and testing conditions in hardware. This approach requires hardware-software co-design and primarily focuses on scalar code. VEGETA [14] presents a set of ISA and microarchitecture extensions over dense matrix engines to support flexible structured sparsity for CPUs. It enables programmable support for diverse deep learning models with varying degrees of sparsity.

VII. CONCLUSION

In this paper, we propose *MACO*, an innovative loosely-coupled multi-core general-purpose architecture specifically designed for GEMM and related applications. In addition, to enhance *MACO*'s programmability and flexibility, we further propose a tile-based instruction set *MPAIS*. Moreover, an efficient mapping scheme for GEMM⁺ workloads on *MACO* is proposed to exhibit the parallel computing power of *MACO*. The aforementioned contributions collectively serve as a promising foundation for future architectural advancements in the realm of GEMM-enhanced general-purpose processors.

ACKNOWLEDGMENT

This work was supported by and National Natural Science Foundation of China (Grant No.U23A20301, No.62202481), TDRCH program (Grant No.22-TDRCH-02-006), KJWPDL program (Grant No.2022-KJWPDL-009), and NUDT Foundation (No.ZK2023-16). Bingcai Sui and Junzhong Shen are co-first authors of this paper. Wei Guo is the corresponding author of this paper.

REFERENCES

- [1] N. P. J. et al., "In-datacenter performance analysis of a tensor processing unit," *ACM SIGARCH Computer Architecture News*, vol. 45, 2017.
- [2] Intel Intel architecture instruction set extensions programming reference,, Std., 2020. [Online]. Available: <https://www.intel.com/doc/accelerator-engines/advanced-matrix-extensions/overview.html>
- [3] G. J. et al., "Rasa: Efficient register-aware systolic array matrix engine for cpu," vol. 2021-December, 2021.
- [4] H. G. et al., "Gemmini: Enabling systematic deep-learning architecture evaluation via full-stack integration," vol. 2021-December, 2021.
- [5] C. L. et al., "Ai accelerator on ibm telum processor industrial product*," 2022.
- [6] H. T. Kung and C. E. Leiserson, "Systolic arrays (for vlsi)," in *Sparse Matrix Proceedings 1978*, vol. 1. Society for industrial and applied mathematics Philadelphia, PA, USA, 1979, pp. 256–282.
- [7] HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers, Std., 2018. [Online]. Available: <https://www.netlib.org/benchmark/hpl/>
- [8] K. e. a. He, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [10] T. e. a. Brown, "Language models are few-shot learners," *NIPS*, vol. 33, pp. 1877–1901, 2020.
- [11] B. e. a. Akin, "Zcomp: Reducing dnn cross-layer memory footprint using vector extensions," in *MICRO*, 2019, pp. 126–138.
- [12] Z. e. a. Gong, "Save: Sparsity-aware vector engine for accelerating dnn training and inference on cpus," in *MICRO*. IEEE, 2020, pp. 796–810.
- [13] S. e. a. Sen, "Sparce: Sparsity aware general-purpose core extensions to accelerate deep neural networks," *IEEE Transactions on Computers*, vol. 68, no. 6, pp. 912–925, 2018.
- [14] G. e. a. Jeong, "Vegeta: Vertically-integrated extensions for sparse/dense gemm tile acceleration on cpus," in *HPCA*. IEEE, 2023, pp. 259–272.