

From Buffers to Registers: Unlocking Fine-Grained FlashAttention with Hybrid-Bonded 3D NPU Co-Design

Jinxin Yu^{1,2,*}, Yudong Pan^{1,2,*}, Mengdi Wang^{1,2}, Huawei Li^{1,2}, Yinhe Han^{1,2}, Xiaowei Li^{1,2}, Ying Wang^{1,2}

¹SKLP, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

{yujinxin21s, panyudong23s, Wangmengdi, lihuawei, yinhes, lxw, wangying2009}@ict.ac.cn

Abstract—Transformer-based models dominate modern AI workloads but exacerbate memory bottlenecks due to their quadratic attention complexity and ever-growing model sizes. Existing accelerators, such as Groq and Cerebras, mitigate off-chip traffic with large on-chip caches, while algorithmic innovations such as FlashAttention fuse operators to avoid materializing large attention matrices. However, as off-chip traffic decreases, our measurements show that on-chip SRAM accesses account for over 60% of energy in long-sequence workloads, making cache access the new bottleneck. We propose 3D-Flow, a hybrid-bonded, 3D-stacked spatial accelerator that enables register-to-register communication across vertically partitioned PE tiers. Unlike 2D multi-array architectures limited by NoC-based router-to-router transfers, 3D-Flow leverages sub-10 μm vertical TSVs to sustain cycle-level operator pipelining with minimal overhead. On top of this architecture, we design 3D-FlashAttention, a fine-grained scheduling method that balances latency across tiers, forming a bubble-free vertical dataflow without on-chip SRAM round-trips. Evaluations on Transformer workloads (OPT and QWEN models) show that our 3D spatial accelerator reduces 46–93% energy consumption and achieves 1.4×–7.6× speedups compared to state-of-the-art 2D and 3D designs.

Index Terms—FlashAttention, 3D, Hybrid-Bonding, Hardware-Dataflow Co-design

I. INTRODUCTION

Large-scale AI inference, particularly Transformer-based large language models (LLMs) [1]–[4], has become increasingly memory-intensive, where the cost of data movement dominates computation. Conventional accelerator [5]–[8] designs devote significant effort to off-chip memory optimization through compression, prefetching, and bandwidth scaling, as well as to **operator fusion** techniques that reduce intermediate memory traffic. Among these, **FlashAttention** [9]–[11] represents a breakthrough: by restructuring attention into fused, tile-based operators, it eliminates the need to materialize the large $N \times N$ attention matrix and reduces off-chip DRAM traffic by more than an order of magnitude. When combined with large on-chip SRAM buffers, such fusion significantly alleviates the off-chip bandwidth bottleneck.

However, as shown in Figure 1, our experiments reveal a new challenge: as off-chip memory traffic decreases, **on-chip SRAM accesses emerge as the dominant energy bottleneck**. Prior studies show that an SRAM access consumes 10–20× more energy than a floating-point multiply-add (FMA) [12], and in FlashAttention workloads with long sequences ($N > 2k$),




Fig. 1: Energy breakdown of operator fusion and unfusion with different sequence lengths for OPT.

on-chip buffer accesses account for over 60% of total energy. This shift highlights that merely scaling on-chip capacity cannot sustain energy efficiency for future LLM accelerators.

One promising direction is to reduce on-chip storage reliance altogether, by spatially unrolling logical resources so that multiple fused operators can execute as a pipeline across processing elements (PEs). Inspired by FlashAttention’s fusion principle, we envision a **register-level pipeline** where operators such as $Q \times K^T$, softmax , and $S \times V$ are mapped onto different PEs and directly pass intermediate data without SRAM round-trips. This approach eliminates most on-chip buffer traffic and enables true dataflow-style execution.

Yet, realizing such fine-grained operator pipelining on large PE arrays faces two major obstacles: 1) **Inter-array connectivity**. Original FlashAttention fuses operators through **buffer-based data exchange** in on-chip SRAM, which still incurs high access cost. To minimize on-chip data movement, intermediate results must instead flow **register-to-register** across PEs, enabling cycle-accurate dataflow execution. However, existing 2D multi-array accelerators rely on **router-to-router NoC transfers** for inter-array communication. This indirection adds multi-cycle latency and energy overhead per hop, making it impossible to sustain fine-grained, cycle-level pipelining. The resulting **pipeline bubbles** cause resource underutilization, undermining the efficiency gains of operator fusion. 2) **Scheduling granularity**. The original FlashAttention fuses operators through **coarse-grained buffering**, but does not support register-to-register pipelining. Achieving **bubble-free, cycle-level** execution across operators requires a new fine-grained scheduling strategy.

To address these challenges, we propose **3D-Flow**, a three-dimensional systolic array architecture enabled by **hybrid bonding** [13], [14]. Unlike micro-bump interconnects with

* indicates equal contribution to this work.

40–50 μm pitch, hybrid bonding provides **sub-10 μm ultra-dense TSV** links with one-cycle latency and low energy per bit, enabling vertical PE-to-PE connections at scale. By vertically stacking PE tiers and assigning consecutive operators to adjacent tiers, 3D-Flow achieves cycle-level operator pipelining without requiring large SRAM buffers.

On top of this hardware substrate, we design **3D-FlashAttention**, a fine-grained scheduling strategy that partitions FlashAttention operators across tiers with **latency-balanced mapping**. This co-design ensures that each tier sustains comparable throughput, forming a tightly coupled vertical pipeline where intermediate results flow directly upward through TSVs. The result is a **bubble-free dataflow execution** that reduces both off-chip and on-chip traffic, significantly improving energy efficiency for Transformer inference. In summary, this paper makes the following contributions:

- **Bottleneck analysis.** We demonstrate that FlashAttention on conventional 2D accelerators is limited by on-chip SRAM accesses, which consume over 60% of energy in long-sequence workloads, and by operator load imbalance that reduces PE utilization.
- **3D-Flow architecture.** We propose a hybrid-bonded, 3D-stacked systolic array where vertically partitioned PE tiers are connected with dense TSVs, supporting register-level operator pipelining with sub-cycle latency.
- **3D-FlashAttention scheduling.** We introduce a fine-grained scheduling method that aligns operator latency across tiers, enabling bubble-free vertical pipelines and eliminating costly SRAM round-trips, thereby sustaining high throughput and energy efficiency for LLM inference. The co-designed Hybrid-bonded NPU architecture can also be generalizable to other fused operators beyond attention.

II. BACKGROUND AND MOTIVATION

A. FlashAttention: From GPU to Systolic Array

FlashAttention [9]–[11] was developed as an IO-aware algorithm that minimizes data movement between high-latency off-chip DRAM and fast on-chip SRAM. It achieves this through two key techniques: tiling and operator fusion, as shown in algorithm 1. By partitioning the query, key, and value matrices into blocks that fit in SRAM, and fusing the QK^T , softmax , and PV steps into a single kernel, FlashAttention avoids writing large intermediate tensors back to DRAM and achieving exceptional performance on NVIDIA A100 and H100 GPUs. However, translating these gains to accelerators built on systolic arrays presents significant challenges that offer considerable optimization potential.

A major bottleneck arises from load imbalance between heterogeneous compute units: QK^T multiplication fully utilizes high-throughput systolic arrays, whereas the subsequent softmax runs on slower scalar or vector units, causing stalls and underutilization of matrix hardware. In addition, on-chip memory bandwidth is constrained, as data transfer between large caches and systolic arrays is serialized over multiple cycles, unlike the parallel shared-memory access of GPUs. This high-latency transfer scales with cache size, making memory access a dominant cost. Overcoming these limitations requires a co-designed architecture and dataflow.

Algorithm 1: FlashAttention-2 forward pass

Input: Matrices $Q, K, V \in \mathbb{R}^{\text{LEN} \times d}$

- 1 Divide Q into $T_r = \lceil \frac{\text{LEN}}{B_r} \rceil$ blocks of size $B_r \times d$ each, and divide K and V into $T_c = \lceil \frac{\text{LEN}}{B_c} \rceil$ blocks of size $B_c \times d$ each;
- 2 **for** $1 \leq i \leq T_r$ **do**
- 3 Initialize $old_m, old_l = (-\infty), (0) \in \mathbb{R}^{B_r}$;
- 4 Initialize $old_O = (0) \in \mathbb{R}^{B_r \times d}$;
- 5 **for** $1 \leq j \leq T_c$ **do**
- 6 $S \leftarrow Q_i K_j^T \in \mathbb{R}^{B_r \times B_c}$;
- 7 $local_m \leftarrow \text{rowmax}(S) \in \mathbb{R}^{B_r}$;
- 8 $new_m \leftarrow \max(local_m, old_m) \in \mathbb{R}^{B_r}$;
- 9 $a \leftarrow old_m - new_m \in \mathbb{R}^{B_r}$;
- 10 $b \leftarrow \exp(\frac{a}{\sqrt{d}}) = \exp 2(\frac{\log_2 e}{\sqrt{d}} a) \in \mathbb{R}^{B_r}$;
- 11 $N \leftarrow S - new_m \in \mathbb{R}^{B_r \times B_c}$;
- 12 $P \leftarrow \exp(\frac{N}{\sqrt{d}}) = \exp 2(\frac{\log_2 e}{\sqrt{d}} N) \in \mathbb{R}^{B_r \times B_c}$;
- 13 $local_l \leftarrow \text{rowsum}(P) \in \mathbb{R}^{B_r}$;
- 14 $new_l \leftarrow old_l \times b + local_l \in \mathbb{R}^{B_r}$;
- 15 $local_O \leftarrow PV_j \in \mathbb{R}^{B_r \times d}$;
- 16 $new_O \leftarrow \text{diag}(b)old_O + local_O \in \mathbb{R}^{B_r \times d}$;
- 17 $old_m \leftarrow new_m$;
- 18 $old_l \leftarrow new_l$;
- 19 $old_O \leftarrow new_O$;
- 20 **end for**
- 21 $O_i \leftarrow \text{diag}(old_l)^{-1}old_O \in \mathbb{R}^{B_r \times d}$;
- 22 **end for**

Several prior works have carried out early-stage investigations into 2D PE arrays. COSA [15] maps FlashAttention onto two collaborating systolic arrays, yet the softmax computation still relies on a dedicated special function unit (SFU). This unit may not deliver enough FLOPs/s to match the throughput of the systolic arrays. FuseMax [16] was the first to integrate FlashAttention entirely within a single systolic array. It implements the rowmax and rowsum steps as spatial reduction operations and concurrently pipelines multiple FlashAttention iterations on the same array. Nevertheless, to enable context switching between iterations, each PE must store intermediate values in 10 registers, resulting in extra register usage. Despite these efforts, the fundamental limitations of 2D PE arrays remain unresolved. In FlashAttention, consecutive operators still rely on the on-chip SRAM for intermediate result exchange, incurring significant and unnecessary memory access overhead. Furthermore, load imbalance among consecutive operators continues to cause underutilization of computational resources and processing stalls.

B. Opportunities Enabled by 3D Integration

Recent advances in semiconductor packaging, exemplified by hybrid bonding, offer a promising path to address the intrinsic dataflow limitations of conventional 2D integration. Traditional planar interconnects suffer from limited wiring density, high latency, and excessive power when scaling to data-intensive workloads. In contrast, 3D integration enables vertical stacking of multiple dies, shortening interconnect distances and increasing inter-die bandwidth density. Hybrid bonding [13], [14] directly connects copper pads at the wafer or die level, achieving ultra-fine-pitch ($< 10 \mu\text{m}$) vertical links with lower resistance and capacitance than TSV- or micro-bump-based approaches [17]. These low-parasitic interconnects reduce transmission delay, raise signaling rates, and improve energy efficiency. The benefits of this packaging paradigm have




Fig. 2: Overview of 3D-stacked PE array architecture and the operator mapping of each layer.

already been realized in commercial products, including high-bandwidth memory (HBM) stacks [18] and 3D SRAM caches such as AMD’s 3D V-Cache [19].

Beyond memory-on-logic stacking, 3D integration allows vertically partitioning the compute fabric itself. Mapping consecutive operators, such as attention sub-stages, onto densely interconnected PE tiers via hybrid-bonded links enables intermediate results to be transferred directly between layers without traversing on-chip SRAM. This vertically stacked architecture reduces on-chip SRAM traffic, and supports operator-specific resource allocation across tiers, thereby alleviating operator load imbalance and sustaining high throughput and energy efficiency for memory-intensive workloads.

C. Motivation

Overall, there are two factors that lead to the inefficiency of running FlashAttention on a 2D systolic array. **First**, due to the physical constraints of the PE array, intermediate results between consecutive operators must be exchanged through the on-chip cache, which creates a bottleneck in on-chip memory accesses. **Second**, load imbalance between consecutive operators frequently leads to idle cycles and underutilization of compute resources. These two factors hinder the efficient execution of FlashAttention on existing 2D PE array based accelerators, which in turn motivates us to explore emerging 3D integration technologies. Through a co-design of architecture and dataflow, we can vertically stack PE arrays with different functionalities, and carefully orchestrate the workload so that consecutive operators are evenly grouped according to load balance and mapped onto the 3D array. In this way, intermediate results can be directly transferred to the next operator via fine-pitch TSV-link enabled by hybrid bonding, significantly eliminating the on-chip memory bottleneck and improving resource utilization, and thereby unlocks the full performance potential of FlashAttention in hardware accelerators.

III. 3D-FLOW ARCHITECTURE

A. Overview

Figure 2 illustrates the overview of the proposed 3D-stacked PE array architecture. Adjacent compute tiers are integrated in a face-to-back configuration, where the registers of spatially aligned PEs are directly interconnected via hybrid-bonded TSV links to enable low-latency data forwarding. Layers with different colors are responsible for distinct operators in FlashAttention, corresponding to the code lines with the same background

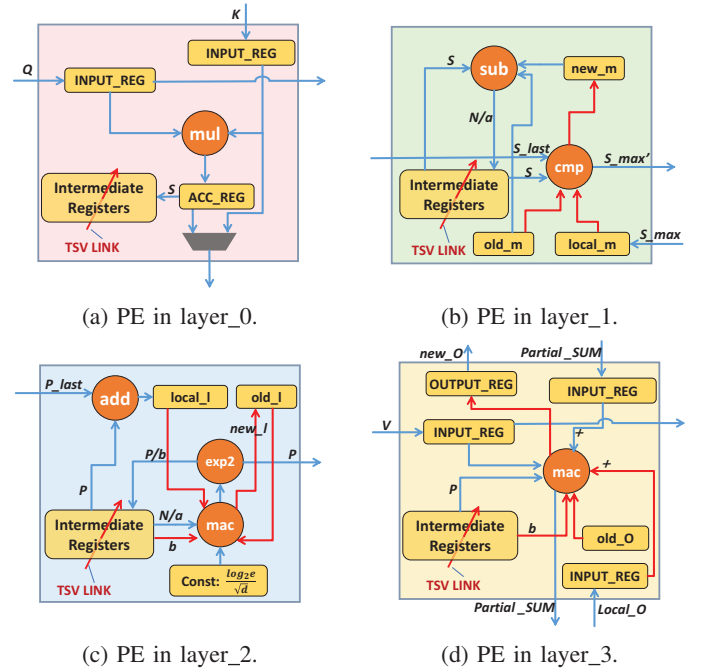


Fig. 3: Custom Design of PE Units in Each Layer.

colors in Algorithm 1. The pipeline of operator flow is carefully orchestrated to ensure balanced workloads across layers. The PE units structure in each layer are custom-designed according to the functionality of the corresponding operators, as shown in Figure 3. A bottom-to-top computation flow corresponds to one iteration of the inner loop in the Algorithm 1, enabling multiple iterations to be executed in parallel through the systolic movement of data within the 3D array. Each PE array layer is configured with dimensions of $d \times d$, where d represents the attention head dimension, thereby allowing a each layer to complete the computation of a whole tile.

B. PE Design

Figure 3 illustrates the detailed microarchitecture of the PE units, each custom-designed for the operators executed in its respective layer. The intermediate registers of corresponding PEs in adjacent tiers are interconnected via TSV links, enabling single-cycle forwarding of intermediate results and avoiding the high-latency SRAM routing in conventional 2D designs. As illustrated in Figure 3, PE in Layer 0 integrates MAC units and is configured with an Outer-Stationary (OS) dataflow to execute the QK^T computation. PE in Layer 1 integrates comparator modules and adders (configured for subtraction) to determine row-wise maximum values and perform subtraction operations, supporting operators such as *RowMax*. PE in Layer 2 integrates exponential computation units, multipliers, and adders, primarily responsible for operators including *exp* and *RowSum*. PE in Layer 3 integrates MAC units along with a reverse data path to implement the *PV* operation and the scaling of the output results. The detailed execution procedures of each operator within the PEs, along with the corresponding dataflow, will be detailed in Section IV-B.

C. Thermal Feasibility

We estimated the thermal behavior of stacked 128×128 PE array. Based on synthesis results, the single PE array has

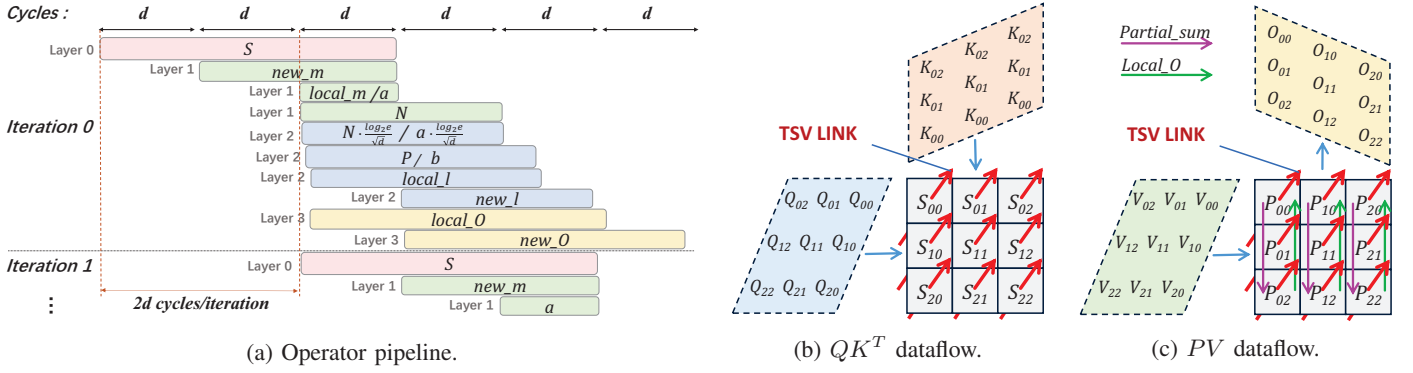


Fig. 4: (a) 3D dataflow and pipeline of consecutive FlashAttention operators across 3D PE array. Operator blocks with different colors correspond to the rows with the same background in Algorithm 1. Each iteration corresponds to one inner loop execution of FlashAttention. d denotes the number of rows/columns in each PE array. When the pipeline is fully filled, one iteration can be completed in $2d$ cycles. (b) Input dataflow of 3D array. (c) Output dataflow of 3D array.

an estimated area of about $A \approx 80 \text{ mm}^2$. At peak activity ($P_{PE} = 200 \text{ } \mu\text{W}$), the power of a PE array is $P_{\text{layer}} \approx 3.3 \text{ W}$, and the four-layer stack consumes $P_{\text{total}} \approx 13.1 \text{ W}$, yielding a layer power density of $\rho \approx 41 \text{ W/cm}^2$. Using a first-order thermal resistance model, the internal temperature rise is small ($\sim 2.8^\circ\text{C}$), indicating good vertical heat conduction. Under a conservative assumption of package thermal resistance $R_{\theta JA} \approx 2.5 \text{ K/W}$ [20], the junction temperature at 25°C ambient is approximately $T_j \approx 83^\circ\text{C}$. With standard advanced packaging techniques, such as thermal TSVs [21], $R_{\theta JA}$ can be further reduced, keeping peak temperatures well within safe operating limits.

IV. 3D-FLASHATTENTION DATAFLOW

A. Overview

As shown in Figure 4a, the consecutive FlashAttention operators are pipelined across different layers of the 3D-stacked PE array. Each block represents a distinct operator, and blocks with the same color are executed in the same layer. The colors of the operators correspond to the background colors of different rows in Algorithm 1. Each iteration corresponds to one inner loop in FlashAttention: given an input $d \times d$ QKV_{tile} , a $d \times d$ attention result is produced, where d denotes both the row/column size of the PE array, which is also equal to the dimension of an attention head. Through careful scheduling, operators are mapped to different layers based on their functionality and computational load, and co-designed with the PE microarchitecture to minimize inter-operator bubbles. When the pipeline is fully filled, only $2d$ cycles are required to complete an iteration, whereas a conventional 2D PE array needs $3d$ cycles merely to perform the QK^T operation. The following section details the execution of the operator flow on the 3D PE array.

B. Layer-wise Dataflow

1) **Layer 0: QK^T :** As shown in Figure 4b, *Layer_0* executes the QK^T operator by feeding Q_{tile} and K_{tile} into the systolic array from horizontal and vertical directions in a parallelogram injection pattern for MAC operations. Partial sums are accumulated within PEs, and after d cycles, $S_{(0,0)}$ is completed and forwarded via TSV to the corresponding PE in *Layer_1*. Thereafter, new S element is produced per cycle,

with all results ready after $3d$ cycles. Meanwhile, as illustrated in Figure 4a, the next iteration of the QK^T operator can start after $2d$ cycles, once the top-left PE becomes idle.

2) **Layer 1: Row-wise Maximum and Subtraction:** The primary function of this layer is to compute the row-wise maximum S_{max} of each row in the matrix S and subtract it. As shown in Figure 4a, after d cycles, the top-left PE receives the first $S_{(0,0)}$ via the TSV_link. Thereafter, in each cycle it propagates rightward, as shown in Figure 3b, where it is compared with the newly received S in the current PE to obtain a larger value S'_{max} , and then continued to be propagated rightward. At $2d$ cycles, the S_{max} of the first row is obtained. Then S_{max} is propagated leftward in the first row, compared with old_m , and the larger value is updated to new_m . Subtraction with both S and old_m yields N and a , respectively. Once N and a are computed, they are transmitted via the TSV_link to the corresponding PE in Layer 2 for the execution of the next operator. Since a exists only in the rightmost PE column, the vector a and the matrix N are completed at $3d$ and $4d$ cycles, respectively.

3) **Layer 2: Exponential-related Operations:** This layer is mainly responsible for the element-wise exponential computations in the *softmax* operation, as well as multiply-accumulate operations involving the scale variable. As shown in Figure 4a, at $2d$ cycles, the first value of the matrix N has been computed and is immediately transmitted via the TSV_link to this layer for subsequent processing. For the exponential function, we employ the more efficient $\exp 2$ implementation commonly used in hardware accelerators [22]. As illustrated in Figure 3c, the constant $\frac{\log_2 e}{\sqrt{d}}$ is stored in the register of each PE and first multiplied with N/a , and the result is then passed to the exponential unit to obtain P/b . Afterwards, P is sent in one direction to the intermediate register and then forwarded via the TSV_link to the corresponding PE in Layer 3, and in the other direction it propagates leftward for accumulation with P from the neighboring PE to produce $local_l$. Finally, $local_l$ undergoes a multiply-accumulate with b and old_l to generate new_l , after which old_l is updated. This layer contains multiple operators that can be executed in parallel, starting from the $2d$ cycle and completing before the $5d$ cycle.

4) **Layer 3: PV and O Scaling:** This layer is primarily responsible for executing the PV operator and scaling the

current output $local_O$ using the previous iteration’s result old_O in the PE’s local registers. As shown in 4a, at the $2d$ cycle, V_tile is fed into the PE array from the left following a parallelogram pattern, and is multiplied with the matrix P received via the TSV_link under a *Weight-Stationary* (WS) dataflow. Partial sums are propagated downward along the vertical direction for accumulation. Once the last row produces $local_O$, the data is transmitted upward to the first row and undergoes a multiply-accumulate operation with old_O and d to generate new_O . The final results are then output from the top of the array. In this layer, the PV computation starts at the $2d$ cycle, and the first $local_O$ is obtained at the $3d$ cycle, after which the O scaling operation is initiated, and all computations are completed at the $5d$ cycle.

Through the aforementioned bubble-free vertical pipelines, a complete inner-loop operator flow of FlashAttention can be completed every $2d$ cycles. Intermediate results between operators are directly transferred via *TSV_links* between registers, thereby avoiding costly SRAM round-trips. The iteration counts of the inner and outer loops in Algorithm 1 are determined by the ratio between the attention sequence length and d . Upon completion of the inner loop, an additional operation is performed, as shown in line 21 of Algorithm 1, to scale O using old_l . When the outer loop finishes, the computation of a single attention head is complete. Multiple heads can be processed in parallel by integrating multiple 3D-stacked PE arrays, depending on the target application requirements and hardware budget.

V. EVALUATION

A. Experiment Setup

Accelerator modeling methodology. We evaluate performance and energy of our 3D architecture and baselines using a cycle-accurate in-house simulator, which integrates Ramulator [23] to capture DRAM access latency and bandwidth, and a custom systolic array simulator to model the specific dataflow we designed for the four-tier architecture. A double-buffering mechanism between memory levels is included to hide access latency. Energy is estimated by tracking hardware activity and invoking Accelergy [24]. To ensure design feasibility, we implemented the complete four-tier PE array in RTL and synthesized it with Synopsys Design Compiler [25] on TSMC 16nm, verifying the functional correctness of our 3D-FlashAttention dataflow and calibrating the energy model. For inter-tier register-to-register communication overhead, we conservatively set the Z-axis data transfer energy [26] to 1.35 pJ/byte, based on stacked DRAM analysis [27], providing a reasonable upper-bound for on-chip register transfers.

Baselines. To comprehensively evaluate the advantages of our proposed 3D architecture and dataflow, we compare it against a series of 2D and 3D baselines. **(1) 2D-Unfused:** A conventional systolic array that executes the phases of attention computation sequentially without operator fusion. **(2) 2D-Fused:** Representative designs, such as FuseMax [16], FLAT [28], and TileFlow [29], which perform deep operator fusion on a single array. **(3) Dual-SA:** A 2D dual-array architecture inspired by COSA [15]. It partitions the QK^T and PV computations onto two arrays connected by a dedicated Softmax unit, enabling intermediate results to flow directly

between these compute units while bypassing on-chip SRAM. **(4) 3D-Base:** An accelerator architecturally identical to ours but maps each operator onto the stacked tiers following the methods in [30], [31], with intermediate results exchanged via SRAM. To ensure a fair comparison, all accelerators are configured with equivalent compute and storage resources (detailed in Table I) and are provisioned with sufficient DRAM. For baselines that require a mapping search (i.e., 2D-Unfused and 3D-Base), we developed a cost model based on Timeloop [32] and utilized linear pruned search strategy to find optimal tiling schemes.

TABLE I: Accelerator Specification in Evaluation

	Ours 3D-Base	2D-Unfused 2D-Fused	Dual-SA
Array Size	$128 \times 128 \times 4$	128×128	$128 \times 128 \times 2$
Clusters	1	4	2
On-Chip Mem. Size	60MB	60MB	60MB
On-Chip BW	8 TB/Sec	8 TB/Sec	8 TB/Sec
Off-Chip BW	400GB/Sec	400GB/Sec	400GB/Sec

Benchmarks. We evaluate the accelerators using representative Transformer models to cover different attention mechanisms: OPT [33] for Multi-Head Attention (MHA) and Qwen [34] for Grouped-Query Attention (GQA). The evaluation spans sequence lengths from 1K to 64K, enabled by implementing dynamic RoPE scaling [35] to extend the models’ context windows beyond their pre-trained limits.

B. Experimental Results

1) Energy consumption: Figure 5 compares the energy consumption of our design against several baselines for the attention computation, with all results normalized to the 2D-Unfused baseline. Overall, our design demonstrates significant energy efficiency, reducing energy by 80.5% to 93% compared to 2D-Unfused approach. Against advanced 2D fusion schemes (FuseMax and Dual-SA), it saves 54.2% to 66.7% in energy. Even compared to the architecturally identical 3D-Naive baseline, our specialized dataflow provides an average energy advantage of 46.8%. This advantage stems from our design’s significant reduction in both costly off-chip DRAM and the increasingly critical on-chip SRAM accesses.




Fig. 5: Comparison of energy consumption of attention with sequence lengths ranging from 1K to 64K.

Figure 6 breaks down data movement of each baseline and reveals the sources of energy savings. Compared to 2D-Unfused, operator fusion methods like FuseMax reduce DRAM accesses by avoiding the materialization of the large intermediate attention score matrix, quadratic in sequence length. However, this shifts the bottleneck to on-chip memory; while FuseMax cuts DRAM accesses by 85.5%, it increases on-chip

SRAM traffic by $2.1\times$, creating a new bottleneck. The Dual-SA architecture attempts to alleviate this with its dual-array design, but its dedicated Softmax unit still relies on SRAM for data exchange, preventing a seamless dataflow between PE arrays. Meanwhile, the 3D-Base baseline leverages 3D stacking for some input reuse (e.g., broadcasting Q-tiles across tiers) but, without a co-designed dataflow, still requires SRAM as an intermediary between operators, leaving 3D integration potential untapped.




Fig. 6: Comparison of average memory DM(data movement volume) of attention.

In contrast, our design leverages 3D-stacked architecture with a co-designed, specialized 3D-FlashAttention dataflow that maps different computational stages to distinct vertical PE array tiers. This strategy not only reduces DRAM accesses via fusion but, more critically, enables intermediate results to flow directly between tiers via registers, completely bypassing on-chip SRAM. Our design reduces on-chip SRAM traffic by an average of 76.6% compared to the other fusion baselines and realizing a truly compute-centric dataflow. Since attention dominates end-to-end energy in long sequences [9], our design yields substantial inference savings, reducing overall energy by 32.7% to 64.2% on average compared to baselines.

TABLE II: Average Energy Breakdown of Our Design

Component	MAC	Reg	SRAM	DRAM	3D-IC Overhead
seq_len=1K	8.5%	21.2%	38.3%	26.7%	5.3%
seq_len=4K	11.7%	31.9%	35.0%	15.1%	6.3%
seq_len=16K	10.4%	29.2%	29.5%	20.8%	10.1%
seq_len=64K	12.0%	34.4%	28.5%	16.2%	8.9%

2) **Energy breakdown and overhead:** We conduct detailed modeling of the energy breakdown for our 3D architecture across various sequence lengths. As shown in Table II, the energy consumption is primarily driven by memory access at different levels and is relatively evenly distributed. Our design incurs additional overhead from increased register access and an average energy overhead of 7.81% due to the use of 3D-IC technology. However, these overheads remain well within acceptable limits.

3) **Speedup and PE utilization:** Figure 7 shows inference speedup of various accelerators, normalized to the 2D-UnFused baseline. Compared to 2D-Unfused, FuseMax, Dual-SA, and 3D-Base baselines, our design achieves average speedups of $7.62\times$, $1.46\times$, $2.36\times$, and $1.43\times$, respectively. This improvement stems from our unique 3D architecture and dataflow co-design, which fundamentally eliminates the latency bottlenecks present in other architectures. Specifically, we map different operators of attention onto stacked PE array tiers, enabling




Fig. 7: Comparison of speedup of attention with sequence lengths ranging from 1K to 64K.

intermediate results to transfer directly via low-latency inter-tier registers. Furthermore, softmax operations are performed within the PE array, avoiding offloading overhead and preventing a separate unit from becoming a system-wide bottleneck. In contrast, while Dual-SA attempts to bypass SRAM by using two PE arrays for QK^T and PV , but sparse and high-latency inter-array interconnects in its 2D layout create significant “drain-and-inject” overhead, and its separate Softmax unit further obstructs the critical path. This high performance is complemented by excellent hardware efficiency: as shown in Figure 8, our 3D architecture achieves an average PE utilization of 87% by reducing DRAM and SRAM accesses to prevent PE stalls and employing a carefully designed inter-tier pipeline for the entire attention computation across vertical tiers, minimizing pipeline bubbles.




Fig. 8: Average utilization of PE arrays with sequence lengths ranging from 1K to 64K.

VI. CONCLUSION

This paper proposes 3D-Flow, a hybrid-bonded, 3D-stacked spatial accelerator that enables register-to-register communication across vertically stacked PE tiers. Unlike 2D multi-array architectures limited by NoC-based router-to-router transfers, 3D-Flow leverages sub-10 μm vertical TSVs to sustain cycle-level operator pipelining with minimal overhead. On top of this architecture, we design 3D-FlashAttention, a fine-grained scheduling method that balances latency across tiers, forming a bubble-free vertical dataflow without on-chip SRAM round-trips. Evaluations on Transformer workloads show that our 3D spatial accelerator reduces 46–93% energy consumption and achieves $1.4\times$ – $7.6\times$ speedups compared to state-of-the-art 2D and 3D designs.

ACKNOWLEDGMENT

This work was supported by the National Natural Science Foundation of China (NSFC) under grant No.92373206. The corresponding authors are Ying Wang, Huawei Li and Xiaowei Li.

REFERENCES

- [1] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:160025533>
- [2] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," 2020. [Online]. Available: <https://arxiv.org/abs/2005.14165>
- [3] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, and F. Azhar, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [4] A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan *et al.*, "Deepseek-v3 technical report," *arXiv preprint arXiv:2412.19437*, 2024.
- [5] 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, 2016.
- [6] Z. Du, R. Fasthuber, T. Chen, P. Inne, L. Li, T. Luo, X. Feng, Y. Chen, and O. Temam, "Shidiannao: Shifting vision processing closer to the sensor," in *Proceedings of the 42nd annual international symposium on computer architecture*, 2015, pp. 92–104.
- [7] 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.
- [8] M. Gao, X. Yang, J. Pu, M. Horowitz, and C. Kozyrakis, "Tangram: Optimized coarse-grained dataflow for scalable nn accelerators," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 807–820.
- [9] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré, "Flashattention: Fast and memory-efficient exact attention with io-awareness," *Advances in Neural Information Processing Systems*, vol. 35, pp. 16 344–16 359, 2022.
- [10] T. Dao, "Flashattention-2: Faster attention with better parallelism and work partitioning," 2023. [Online]. Available: <https://arxiv.org/abs/2307.08691>
- [11] J. Shah, G. Bikshandi, Y. Zhang, V. Thakkar, P. Ramani, and T. Dao, "Flashattention-3: Fast and accurate attention with asynchrony and low-precision," 2024. [Online]. Available: <https://arxiv.org/abs/2407.08608>
- [12] M. Horowitz, "1.1 computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. IEEE, 2014, pp. 10–14.
- [13] Y. Kagawa, S. Hida, Y. Kobayashi, K. Takahashi, S. Miyanomae, M. Kawamura, H. Kawashima, H. Yamagishi, T. Hirano, K. Tatani *et al.*, "The scaling of cu-cu hybrid bonding for high density 3d chip stacking," in *2019 electron devices technology and manufacturing conference (EDTM)*. IEEE, 2019, pp. 297–299.
- [14] M.-C. Lu, "Hybrid bonding for ultra-high-density interconnect," *Journal of Electronic Packaging*, vol. 146, no. 3, p. 030802, 2024.
- [15] Z. Wang, G. Wang, H. Jiang, N. Xu, and G. He, "Cosa: Co-operative systolic arrays for multi-head attention mechanism in neural network using hybrid data reuse and fusion methodologies," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2023, pp. 1–6.
- [16] N. Nayak, X. Wu, T. O. Odemuyiwa, M. Pellauer, J. S. Emer, and C. W. Fletcher, "Fusemax: Leveraging extended einsums to optimize attention accelerator design," *arXiv preprint arXiv:2406.10491*, 2024.
- [17] J. Wang, F. Duan, Z. Lv, S. Chen, X. Yang, H. Chen, and J. Liu, "A short review of through-silicon via (tsv) interconnects: metrology and analysis," *Applied Sciences*, vol. 13, no. 14, p. 8301, 2023.
- [18] H. Jun, J. Cho, K. Lee, H.-Y. Son, K. Kim, H. Jin, and K. Kim, "Hbm (high bandwidth memory) dram technology and architecture," in *2017 IEEE International Memory Workshop (IMW)*. IEEE, 2017, pp. 1–4.
- [19] J. Wu, R. Agarwal, M. Ciraula, C. Dietz, B. Johnson, D. Johnson, R. Schreiber, R. Swaminathan, W. Walker, and S. Naffziger, "3d v-cache: the implementation of a hybrid-bonded 64mb stacked cache for a 7nm x86-64 cpu," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65. IEEE, 2022, pp. 428–429.
- [20] "Semiconductor and ic package thermal metrics," <https://www.ti.com/lit/an/spra953d/spra953d.pdf>, 2025.
- [21] Z. Wang, R. Dong, R. Ye, S. S. K. Singh, S. Wu, and C. Chen, "A review of thermal performance of 3d stacked chips," *International Journal of Heat and Mass Transfer*, vol. 235, p. 126212, 2024.
- [22] M. A. Hussain, S.-W. Lin, and T.-H. Tsai, "An area-efficient and high throughput hardware implementation of exponent function," in *2022 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2022, pp. 3369–3372.
- [23] H. Luo, Y. C. Tuğrul, F. N. Bostancı, A. Olgun, A. G. Yağlıkçı, and O. Mutlu, "Ramulator 2.0: A modern, modular, and extensible dram simulator," *IEEE Computer Architecture Letters*, vol. 23, no. 1, pp. 112–116, 2023.
- [24] Y. N. Wu, J. S. Emer, and V. Sze, "Accelerger: An architecture-level energy estimation methodology for accelerator designs," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–8.
- [25] Synopsys, Inc., "Design compiler rtl synthesis," <https://www.synopsys.com/implementation-and-signoff/rtl-synthesis-test/dc-ultra.html>, 2025.
- [26] D. Rich, A. Kasperovich, M. Malakoutian, R. M. Radway, S. Hagiwara, T. Yoshikawa, S. Chowdhury, and S. Mitra, "Thermal scaffolding for ultra-dense 3d integrated circuits," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2023, pp. 1–6.
- [27] R. Mathur, A. K. A. Kumar, L. John, and J. P. Kulkarni, "Thermal-aware design space exploration of 3-d systolic ml accelerators," *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits*, vol. 7, no. 1, pp. 70–78, 2021.
- [28] S.-C. Kao, S. Subramanian, G. Agrawal, A. Yazdanbakhsh, and T. Krishna, "Flat: An optimized dataflow for mitigating attention bottlenecks," in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 2023, pp. 295–310.
- [29] S. Zheng, S. Chen, S. Gao, L. Jia, G. Sun, R. Wang, and Y. Liang, "Tileflow: A framework for modeling fusion dataflow via tree-based analysis," in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture*, 2023, pp. 1271–1288.
- [30] J. M. Joseph, A. Samajdar, L. Zhu, R. Leupers, S. K. Lim, T. Pionteck, and T. Krishna, "Architecture, dataflow and physical design implications of 3d-ics for dnn-accelerators," in *2021 22nd International Symposium on Quality Electronic Design (ISQED)*. IEEE, 2021, pp. 60–66.
- [31] H.-T. Kung, B. McDanel, and S. Q. Zhang, "Mapping systolic arrays onto 3d circuit structures: Accelerating convolutional neural network inference," in *2018 IEEE International Workshop on Signal Processing Systems (SiPS)*. IEEE, 2018, pp. 330–336.
- [32] A. Parashar, P. Raina, Y. S. Shao, Y.-H. Chen, V. A. Ying, A. Mukkara, R. Venkatesan, B. Khailany, S. W. Keckler, and J. Emer, "Timeloop: A systematic approach to dnn accelerator evaluation," in *2019 IEEE international symposium on performance analysis of systems and software (ISPASS)*. IEEE, 2019, pp. 304–315.
- [33] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin *et al.*, "Opt: Open pre-trained transformer language models," *arXiv preprint arXiv:2205.01068*, 2022.
- [34] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, "Qwen technical report," *arXiv preprint arXiv:2309.16609*, 2023.
- [35] N. Shang, L. L. Zhang, S. Wang, G. Zhang, G. Lopez, F. Yang, W. Chen, and M. Yang, "Longrope2: Near-lossless llm context window scaling," *arXiv preprint arXiv:2502.20082*, 2025.