

SATA: Sparsity-Aware Scheduling for Selective Token Attention

Zhenkun Fan¹, Zishen Wan¹, Che-Kai Liu¹, Ashwin Sanjay Lele², Win-San Khwa³, Bo Zhang²
Meng-Fan Chang³, Arijit Raychowdhury¹

¹*Electrical and Computer Engineering Department, Georgia Institute of Technology, Atlanta, GA, USA*

²*TSMC Corporate Research, San Jose, CA, USA*

³*TSMC Corporate Research, Hsinchu, Taiwan*

Email: {zfan87, zwan63}@gatech.edu

Abstract—Transformers have become the foundation of numerous state-of-the-art AI models across diverse domains, thanks to their powerful attention mechanism for modeling long-range dependencies. However, the quadratic scaling complexity of attention poses significant challenges for efficient hardware implementation. While techniques such as quantization and pruning help mitigate this issue, *selective token attention* offers a promising alternative by narrowing the attention scope to only the most relevant tokens, reducing computation and filtering out noise.

In this work, we propose SATA, a *locality-centric dynamic scheduling scheme* that proactively manages sparsely distributed access patterns from selective Query-Key operations. By reordering operand flow and exploiting data locality, our approach enables early fetch and retirement of intermediate Query/Key vectors, improving system utilization. We implement and evaluate our token management strategy in a control and compute system, using runtime traces from selective-attention-based models. Experimental results show that our method improves system throughput by up to 1.76× and boosts energy efficiency by 2.94×, while incurring minimal scheduling overhead.

I. INTRODUCTION

Since Transformers revolutionized natural language processing (NLP) with their breakout success across multiple tasks [1], the attention mechanism has become foundational in modern machine learning. Transformers now power state-of-the-art models in computer vision [2]–[5], natural language processing [6]–[9], multimodal learning and agentic tasks [10]–[14], laying the groundwork for the next generation of AI algorithms. Despite their impressive capabilities, Transformers come with high deployment overhead, often requiring multiple GPUs and incurring substantial energy consumption [15].

The power of attention lies in its global receptive field, allowing each token to attend to all others within a context. However, this expressiveness also leads to the well-known quadratic complexity in sequence length, making the attention layer a computational bottleneck. As model sizes grow, attention increasingly becomes memory-bound [16], [17], posing challenges for scaling to edge platforms. Consequently, both academia and industry have been actively developing Transformer acceleration techniques to improve throughput and energy efficiency [18]–[21].

To reduce redundant computation, recent work has explored sparsity in Transformers. These models often exhibit inherent

redundancy in token, head, or weight-level representations. Sparse Transformer accelerators [22]–[25] have leveraged this property to reduce workload and boost efficiency while maintaining accuracy. In particular, TopK Selective Query-Key Attention [26]–[30] has gained traction by focusing on the most influential key tokens for a given query, thereby pruning less important MAC operations in the attention score computation. This algorithmic sparsity opens new opportunities for hardware-software co-design to optimize performance.

At the core of the attention mechanism lies the matrix multiplication (MatMul) primitive. As a fundamental building block of ML workloads, MatMul has seen significant hardware optimization over the past decade. Compute-in-Memory (CIM) architectures are a notable direction, offering latency and energy benefits by minimizing memory movement and enabling massive parallelism. Recent Transformer accelerators have extended CIM capabilities with features like interconnect reconfigurability [31], [32], bitwidth adaptability [33], and algorithm-hardware co-design [34]. Due to its dynamic nature, selective Q-K attention often leads to fragmented operational flow. Harvesting energy savings by halting the corresponding functional unit brings down hardware utilization. In the worst case, the reduced operand reuse distance can induce a surge of external memory access, degrading system efficiency and throughput.

In this work, we propose SATA, a *locality-centric dynamic scheduling scheme* that reorders Query and Key access patterns to improve operand locality in attention computation. Our scheduler targets Scaled Dot-Product Attention (SDPA) across Multi-Head Attention (MHA) layers, executing only the most impactful QK-MAC operations while maximizing system utilization. Through centralizing Multiply and Accumulates (MACs), we minimized the retention duration of Q/K operand without sacrificing model accuracy. Combined with an optimized scheduler architecture, the system utilization is maximized with parallelized Q, K operations

This paper, therefore, makes the following contributions:

- We propose a sparsity-aware scheduling scheme that enhances data reuse in sparse Query-Key attention. By minimizing reuse distance through sorted operand access, the scheme improves throughput and energy efficiency.

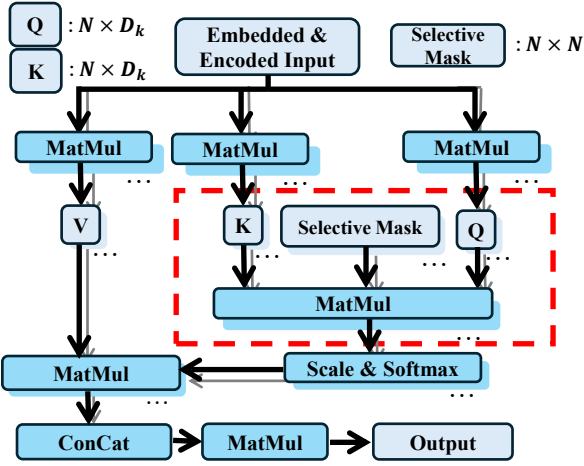


Fig. 1: MHA visualization. Red dashed box represents the targeted workload in SATA. N : sequence length (number of tokens); D_k : embedding dimension of Query and Key.

- We design a lightweight controller to implement the proposed scheduling scheme within digital design flow. Experimental results show the scheduling overhead is just 2.2% in the most energy-sensitive workload, with a worst-case overhead of 5.9%.
- We develop a comprehensive benchmarking framework using a silicon-validated CIM simulator and metadata from commercial EDA tools, evaluating across four Transformer models. Our design achieves throughput gains of $1.47\times$, $1.76\times$, $1.59\times$, and $1.5\times$, and energy efficiency improvements of $1.81\times$, $2.1\times$, $1.85\times$, and $2.94\times$ on TTST, KVT-DeiT-Tiny, KVT-DeiT-Base, and DRFormer, respectively. The code is publicly available at github.com/SenFFF/SATA.

II. BACKGROUND AND RELATED WORK

A. Attention Mechanism

The attention mechanism (Fig. 1) is first introduced in Neural Machine Translation [35]. In a series of tokenized inputs, each token searches through the sequence and generates scores for other tokens. It then forms a feature representation with MAC upon vectors embedded in Value space weighted by the relevance scores. Despite state-of-the-art performance, modern transformers grow exponentially in terms of size and cost. Commercial products rely heavily on datacenters. Transmission between cloud and edge has incurred great overhead and security concerns. Synergistic approaches to slim down transformers are a heated topic in academia and industry.

B. Attention Sparsity

Transformer tokens are projected to high-dimensional spaces before further proceeding. Such redundancy is vital to training. Inference-centric models, on the other hand, have an effective datapath where tokens can be identified across the spectrum from critical to obsolete. Identifying inessential

components and avoiding related compute have great potential for transformer acceleration.

Among all techniques, TopK selective attention is a more flexible scheme where each Query(Q) is computed against only a subset of Keys(K) [26], [27]. As dense transformers are often considered overparameterized, selective attention serves as a regularization method providing robust performance and sheds light on a path for efficient hardware implementation.

C. Transformer Accelerators

Accelerating transformers has been one of the hottest topics in the past decade. For sparsity-driven approaches, A^3 [22] performed successive approximation. SpAtten [23] combined cascaded token pruning, head pruning, TopK and dynamic quantization to alleviate DRAM access. Energon [24] implemented selective attention by progressively filtering trivial keys. Inspired by the dynamic philosophy in [33], SATA performed dynamic scheduling in MHA (Fig. 1) to improve throughput and system efficiency.

III. SCHEDULING SCHEME

In this section, we present the SATA scheduling scheme. We begin by characterizing workload, then describe the algorithmic flow, and finally cover the hardware implementation. The notation for parameters can be found in Algo. 1, 2 and Tab. I.

A. Selective Attention Layer

MHA can be decomposed to projection, Q-K MatMul, A-V MatMul, Feed Forward Network (FFN), and nonlinear operations. Projection and FFN are considered Static MatMul. In contrast, Q-K and A-V MatMuls rely on vectors generated on the fly and are input-dependent. These dynamic MatMul scale non-linearly, hindering attention's edge deployment [16].

We focus therefore on Q-K MatMul across heads. Here, the input to SATA is the TopK indices of Keys relevant to Queries. The acquisition of indices has been studied extensively [19], [23], [24] and its cost is integrated in evaluation (Sec. IV).

B. Intra-Head Mask Sorting

Fig. 2(a) shows how SATA sorts selective masks. Scattered K access often leads to redundant operand fetches. SATA introduces intra-head mask sorting, where Q_s and K_s are reordered to improve operand locality before scheduling.

The process begins by randomly picking a K index as the pointer. Its access pattern (mask column), copied as vector *dummy*, serves as a reference vector throughout sorting. Unsorted K access vectors are compared with *dummy* to generate similarity scores. The K index with the highest score would be marked as the next K pointer, and its access pattern is used to increment *dummy*. This process loops till all Keys are sorted and generates a list of sorted K indices. (Algo. 1, line 4-12) Despite the order of $O(n^2)$, the overhead of sorting is light as the mask is binary. Sorting is further optimized in Sec. III-E.

With sorted K_s , Q_s would be classified into three groups: HEAD, TAIL or GLOB, based on sorted distribution. Classification relies on a dynamic 'Heavy Size' S_h . Specifically,

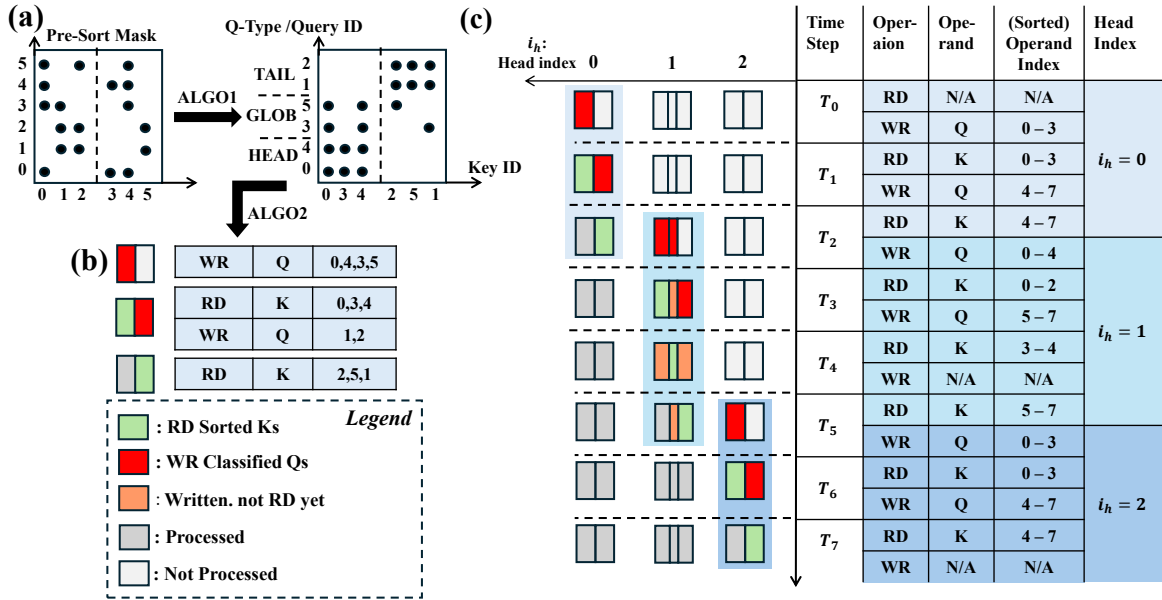


Fig. 2: Visualization for SATA algorithm. (a),(b) present a demo for Algo. 1, 2 ($N = 6$, $S_h = \frac{N}{2}$). Q/K indices are original token indices. This head is classified as HEAD since the number of GLOB Qs is smaller than $\frac{N}{2}$. Tie is broken by assigning head condition to HEAD when number of HEAD-Qs equals that of TAIL-Qs. (c) presents a scheduled demo. Two heads ($i_h = 0, 2$) are perfectly sorted with $S_h = \frac{N}{2}$. The rest ($i_h = 1$) represents a more general case ($S_h < \frac{N}{2}$) where conceding (S_h decrement) has happened.

Algorithm 1 Key Sorting and Query Classifying

```

1: Input: #Token= $N$ , Selective Mask  $QK \in \{0, 1\}^{N \times N}$ , threshold  $\theta$ 
2: Output: Sorted Key-Order  $Kid$ , Q-Types  $QT$ , Head-Type  $HT$ .
3:
4:  $Kid = []$ ,  $Dummy = \mathbf{0}^N$  // Init
5: // Rand Seed
6:  $\_kid = \text{rand}(N)$ ,  $Dummy.update(QK[:, \_kid])$ ,  $Kid.append(\_kid)$ 
7: // Sort
8: for  $i$  in  $0 \dots N - 2$  do
9:    $\_kid \leftarrow \arg \max(Dummy^T \cdot QK[:, i])$ ,  $i \notin Kid$ 
10:   $Kid.append(\_kid)$ 
11:   $Dummy.update(QK[:, \_kid])$ 
12: end for
13: // Classify
14:  $S_h = \frac{N}{2}$ ,  $QK_s = QK[:, Kid]$ ,  $QT = []$ , #Glob = 0
15: while True do
16:   for  $i$  in  $0 \dots N - 1$  do
17:      $QT.update(\text{classify}(QK_s[i, :], S_h))$ 
18:     #Glob += 1 if (i-th Q is GLOB)
19:   end for
20:   if #Glob >  $\theta$  then
21:      $S_h \leftarrow S_h - 1$  // Escape GLOB w/ smaller  $S_h$ 
22:     continue
23:   else
24:      $HT \leftarrow \#(\text{HEAD}) > \#(\text{TAIL}) ? \text{HEAD} : \text{TAIL}$ 
25:     Break
26:   end if
27: end while
28: return  $Kid$ ,  $QT$ ,  $S_h$ ,  $HT$ 

```

- Qs not accessing last S_h sorted Ks are tagged as HEAD.
- Qs not accessing first S_h sorted Ks are tagged as TAIL.
- Other Qs are tagged GLOB which represents less locality.

A head is local as long as GLOB Qs do not dominate. The exact type depends on the dominant Q-type being HEAD or

TAIL (head-type HEAD or TAIL). Otherwise, the head is in a GLOB state. Should that happen, S_h is decremented and Qs are reclassified to escape from GLOB status (Algo. 1, line 14-27). Packed together, this two-step process enhances data reuse and lays the foundation for sparsity aware scheduling.

C. Sparsity-Aware Inter-Head Scheduling

For selective token mechanism, a straightforward approach to reduce energy is to gate compute units when loading operands (e.g., trivial Keys). Yet, such pruning brings marginal benefits because it fails to fully exploit system utilization. SATA goes beyond pruning to sparsity-aware inter-head scheduling, with the central idea of dynamically orchestrating Q-K operations across heads so that system modules (e.g., Systolic array PEs or CIM tiles) remain fully utilized.

Unlike prior works that fix Ks as the stationary operand [31]–[34], SATA designates Qs to be stationary due to their low variance of arithmetic intensity. Specifically, each Q would always have the same amount of K access, while Ks behave otherwise [26], [30].

The scheduling flow can be formalized as a Finite State Machine (FSM). Each state corresponds to a different situation scheduling HEAD/TAIL/GLOB sets:

- **init:** Load major Qs, which mean HEAD&GLOB for condition HEAD, or TAIL&GLOB for condition TAIL.
- **intoHD:** Launch MatMul with the first $[0 : S_h - 1]$ sorted Ks while loading minor Qs (TAIL Qs for condition HEAD, or HEAD Qs for condition TAIL). This is possible since sorting has granted HEAD (TAIL) Qs

Algorithm 2 Sparsity-Aware Scheduling

```

1: Input: #Token= $N$ , #Heads= $N_h$ , head-Types  $HT$ ,
   sorted Key indices for all heads  $Kid, \in \mathbb{N}^{N_h \times N}$ , classified Q-Types
    $QT_i = [Q_{i,HEAD}; Q_{i,TAIL}; Q_{i,GLOB}], i \in [0, \dots, (N_h - 1)]$ , Heavy-
   Size  $S_{h,i}, i \in [0, \dots, (N_h - 1)]$ ;
2:
3: Output: Q-Load Sequence QSeq, K-MAC Sequence KSeq
4: QSeq = [], KSeq = [], Head-id  $i_h = 0$ , state = intoHD // Init
5: while True do
6:   if state is intoHD then
7:     if  $HT_{i_h}$  is HEAD then
8:       QSeq.Update(  $Q_{i_h,HEAD} + Q_{i_h,GLOB}$  )
9:     else if  $HT_{i_h}$  is TAIL then
10:      QSeq.Update(  $Q_{i_h,TAIL} + Q_{i_h,GLOB}$  )
11:     end if
12:     // Finish reading K of head  $i_{h-1}$ 
13:     KSeq.Update(  $Kid[i_h - 1, N - S_h : ]$  )
14:     state  $\leftarrow$  midstHD
15:   else if state is midstHD then
16:     KSeq.Update(  $Kid[i_h, S_h : N - S_h]$  )
17:     state  $\leftarrow$  outtaHD
18:      $i_h \leftarrow i_h + 1$ 
19:   else if state is outtaHD then
20:     if  $HT_{i_h}$  is HEAD then
21:       QSeq.Update(  $Q_{i_h,TAIL}$  )
22:     else if  $HT_{i_h}$  is TAIL then
23:       QSeq.Update(  $Q_{i_h,HEAD}$  )
24:     end if
25:     KSeq.Update(  $Kid[i_h, : S_h]$  )
26:     state  $\leftarrow$  intoHD
27:   end if
28:   if  $i_h == N_h$  then
29:     break
30:   end if
31: end while
32: return QSeq, KSeq

```

would not require the last (first) S_h sorted Ks to generate score matrix, which enabled parallel execution.

- *midstHD*: This state is activated only when S_h is not half the tile size. Ks with sorted indices $[S_h : N - S_h]$ perform MAC with every Q to preserve model accuracy.
- *outtaHD*: Conclude MAC with Ks indexed $[N - S_h : N]$ and start loading major Qs for the next head. Since the current major Qs need not access the leftover Ks, they can be safely retired and release storage capacity.
- *wrapGLOB*: For heads in GLOB state, perform Q-K MatMul in a 'Load Q then MAC with K' manner.

The scheduling stops when all LOCAL heads are consumed. For the remaining GLOB heads, the scheduler reverts to conventional computation flow. In experimentation, most of the GLOB heads can escape GLOB state with a smaller S_h . The profiled data shows $< 0.1\%$ GLOB heads are observed among 2K Traces from TTST.

Algo. 2 summarizes scheduling. Note that edge transition and GLOB wrapping up are omitted for brevity. Sorting and scheduling are visualized in Fig. 2.

D. Scale to Larger Sequence Length

Scaling to longer sequences is one of the critical hurdles in both the algorithmic and hardware communities. A growing sequence length N incurs quadratic growth in Q-K MatMul and extensive off-chip traffic.

SATA introduces tiling within each head to create smaller sub-blocks (e.g., 16×16 tiles). Each tile is executed like a 'subhead'. Sorting would be conducted across Q-folds while fold-wise Ks are reused. A similar process would then repeat across K-folds and complete the score matrix. This brings buffer demand to become manageable.

In the original sorting scheme, every Key and Query is determined indispensable. This is invalidated in tiled subheads. As such, zero-skip is introduced to the tiled selective masks. The scheduler identifies redundant Query(Key) in the tiled matrix by a column(row)-wise reduction AND operation. Together, tiling and zero-skip enable SATA to scale to long sequences with preserved locality, while preventing the quadratic growth to overwhelm hardware.

E. Hardware Implementation

SATA scheduling scheme is implemented with a lightweight scheduler integrated with compute engine. Its functionality is to coordinate memory traffic in and out of the compute unit.

We decompose the functionalities of SATA and organize them into digital modules as in Fig. 3a. Zero-unit detects and filters redundancy. Dot-product engine updates the similarity score and sorts Keys with a priority encoder. Key order is staged in a FIFO during sorting. Query order is written to another FIFO based on the classification result. FIFOs can be read by computational parts for operand retrieval.

The most energy and runtime consuming step for scheduling is dot products (line 7-12 in Algo. 1). Instead of directly computing similarity and updating the dummy vector (Eq. 1), SATA saves the cumulative distance between the *Dummy* vector and all mask columns ($QK[:, i], i \in [0, N - 1]$). Whenever a Key index j is sorted, the Psum Registers are incremented by the binary dot-product result between $QK[:, j]$ and all unsorted columns (Eq. 2).

$$Distance_i = Dummy^T \cdot QK[:, i] \quad (1)$$

$$Psum-Reg[i] += QK[:, i]^T \cdot QK[:, j], \quad i \notin Kseq \quad (2)$$

This essentially eliminates the repetitive MAC operation between sorted Ks and unsorted ones, and delivered design with better PPA metrics. The overhead of the scheduler is discussed in Sec. IV-D.

IV. EVALUATION

This section evaluates SATA's efficiency, scalability, hardware overhead, and compatibility with SOTA accelerators. We first introduce the estimation framework and workloads, then present throughput, energy, and SOTA-compatibility results.

A. Estimation Framework and Testing Workloads

1) *PPA estimation*: We access SATA with an estimation framework, which models both algorithmic behavior and hardware activities. Evaluated system consist of the computational CIM module (Fig. 3c) and the scheduler (Fig. 3a). We implement scheduler in SystemVerilog and synthesized it with TSMC65nm process in Design Compiler. Place and Route is

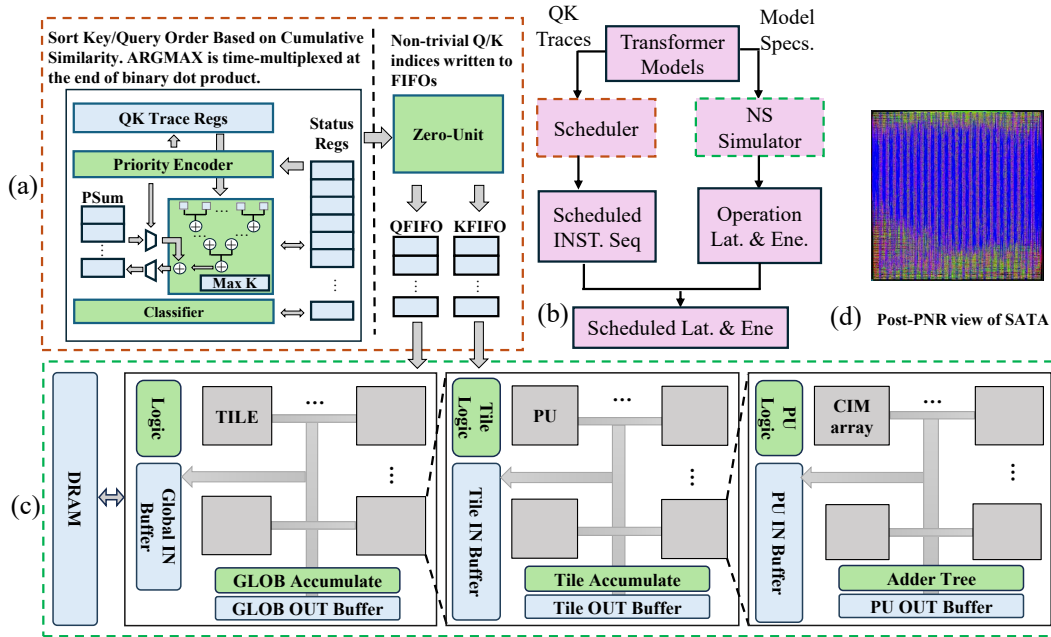


Fig. 3: SATA estimation framework. (a) Schematic view of the proposed scheduler. Status-Regs control the scheduling process and stage intermediate data. (b) Workflow of SATA estimation framework. (c) Homogeneous CIM-centric computational system initialized by NeuroSim. Input activations start from DRAM, execute MAC operation on sub-arrays, and are transferred back to DRAM to finish computation. (d) Post-PNR figure of SATA.

TABLE I: Workload Specification & Post-Schedule Statistics.

| TopK Model | EMB-DIM (D_k) | K#/Token | 0-Skip | Dataset | GlobQ% | Tile Size(S_f) | Avg Heavy-Size (S_h) | Avg #($S_h=1$) |
|--------------------|-------------------|----------|--------|---------|--------|--------------------|--------------------------|------------------|
| TTST [30] | 65536 | 15/30 | 0 | [36] | 24.2% | N | $0.463N$ | 1.55 |
| KVT-DeiT-Tiny [27] | 64 | 50/198 | 1 | [37] | 33.3% | $0.11N$ | $0.053N$ | 0.62 |
| KVT-DeiT-Base [27] | 64 | 64/198 | 1 | [37] | 46.4% | $0.11N$ | $0.051N$ | 1.38 |
| DRSformer [28] | 4800 | 12/48 | 1 | [38] | 14.8% | $0.125N$ | $0.062N$ | 0.05 |

done with IC Compiler2. Both SATA and CIM adopt 1GHz frequency clock rate.

CIM module is estimated with NeuroSim, which is validated and calibrated against post-silicon measurements [39]. The CIM architecture is a multi-level, homogeneous system that takes both caches and interconnect into consideration (Fig. 3c). We supplement the original dense CIM engine with SATA to quantify its throughput and energy benefit.

We size the CIM subarray to 32×32 . The technology is adopting 65 nm process metadata. For additional system configuration details, we redirect interested readers to [40]. We performed Design Space Exploration (DSE) on the SATA configuration to ensure optimal performance is delivered.

The estimation workflow (Fig. 3b) starts with trace extraction. Scheduler produces operator with QK traces. The instructions (RD/MAC, WR/Load) are combined with a customized NeuroSim codebase to estimate latency and energy.

a) Throughput: The throughput benefit of SATA arises from reduced modules idleness after interleaving Q/K operations. We extract metadata of data transmission (DT) and Compute (Comp) latency from NeuroSim. The latency(τ) for a scheduled time step that reads (MACs) x Ks and writes (loads) y Qs is estimated as

$$\tau_i = \min(\tau_{RD,DT} \cdot x, \tau_{WR,ARR} \cdot y) + \min(\tau_{RD,COMP} \cdot x, \tau_{WR,DT} \cdot y) \quad (3)$$

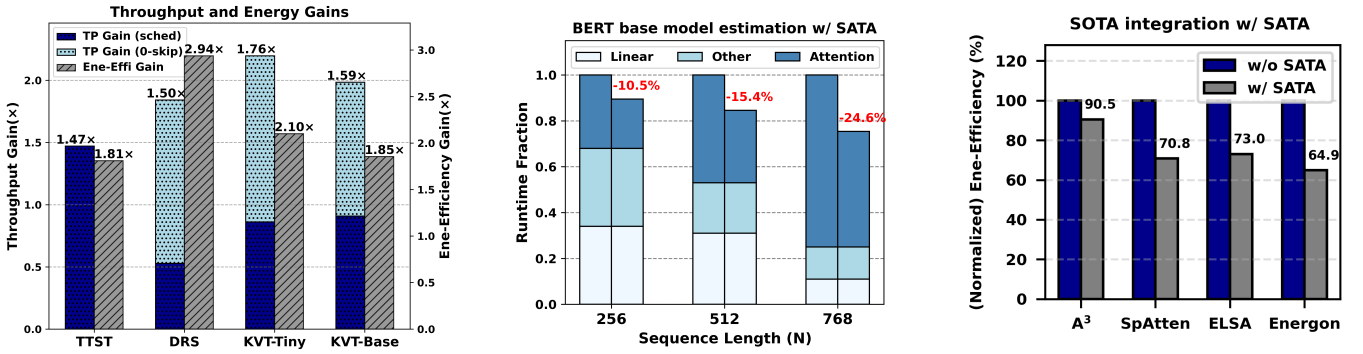
The latency of scheduling is negligible compared to compute ($< 5\%$) and can be hidden through pipelining.

b) Energy Estimation: NeuroSim estimates energy assuming dense MatMul. In scheduled operations, only a subset of tiles perform MAC, while the remaining can update weights (Qs). Assume a head is sorted to have a HEAD Qs, b TAIL Qs, and c GLOB Qs ($S_h < \frac{N}{2}$). Then Keys indexed $[N - S_h : N]$ would bypass a Qs; Keys indexed $[0 : S_h]$ would bypass b Qs. MACs are assumed to be conducted in a dense manner, albeit in a subset of tiles. The energy of SATA is extracted from EDA reports.

2) Selective Transformer Models: We benchmark SATA against k-NN Attention [27], TTST [30], and DRSformer [28], all of which adopted TopK attention. Model parameters are listed in Tab. I. Values of K are chosen, such that the model performance degradation is negligible.

B. Performance and Algorithmic Parameter Analysis

Fig. 4a shows the scheduled Q-K attention performance result for TTST. The throughput benefit arises from paral-



(a) QK throughput and energy efficiency gain of SATA. The cost of QK index compute and scheduler has been incorporated.

(b) Normalized BERT based model estimation with SATA integration [24].

(c) Energy-efficiency gain by integrating SATA into SOTA accelerators.

Fig. 4: SATA evaluation results for (a) QK throughput, energy efficiency gain for different workloads; (b) Self-Attention runtime reduction with SATA; and (c) SOTA energy-efficiency improvement with SATA.

lized scheduled operation, hence improved hardware utilization, while the energy reduction comes from pruned MAC operations. We initialize $S_h = \frac{N}{2}$ and $\theta = \frac{N}{2}$. $S_h = \frac{N}{2}$ represents the optimistic case, though relaxation often happens to escape **GLOB** state. In practice, the average post-schedule S_h stays close to the optimal condition (50% tile size), as reported in Tab. I. The average S_h decrement times are also given in Tab. I. This incurs minimal overhead, as the latency and energy consumption of classification is minor in scheduler.

We also integrate the QK-index computation and SATA scheduling cost into estimation shown in Fig. 4a [23], [24]. SATA achieves a peak throughput gain of 1.76 \times and an energy efficiency benefit of 2.94 \times . Moreover, SATA is compatible with MatMul engine beyond CIM. A preliminary TTST test on a SATA-enhanced systolic array platform, modeled with ScaleSIM [41], demonstrated a 3.09 \times throughput gain, with stall cycles reduced from 90.4% to 75.2%.

C. Scaling capability

In SATA, long vectors are difficult to sort algorithmically, and the scheduler itself becomes prohibitive due to the logarithmic growth of tree structure. We pick tile size (S_f) such that SATA delivers the highest performance. The S_f values for different workloads are recorded in Tab. I.

As S_f decreases, throughput gain initially increases as SATA enables higher system utilization. When S_f becomes even smaller, the benefit from zero-skip starts to dominate (exceeding 50%). Redundant operands will not be pushed into FIFOs, which means higher fraction of trivial operands inherently makes scheduling contributions less significant.

D. Scheduler Overhead

We evaluate the overhead of SATA scheduler compared to computational module. The cost is negligible with large embedding dimensions but scales quadratically with tile size (register array) and logarithmically with tree-style modules. Scheduling latency can be effectively hidden through pipelining as long as scheduler takes less time than Q-K MatMul.

Generally, latency overhead remains minor ($< 5\%$) when $D_k \geq 64$, or when $S_f \leq 24$, compared to an optimized CIM core. In terms of energy, a minor cost ($< 5\%$) assumption fails when $D_k < 32$ or $S_f > 28$, which already falls out of the optimal gain setting reported in Tab. I.

E. Integrating SATA to Existing Accelerators

Previous accelerators [19], [22]–[24] execute sparse Q-K MAC after index acquisition. Though parallel hardware or pipelined designs improve efficiency, their sparsified nature remain underutilized. By integrating SATA, additional energy and throughput benefits can be achieved through localized operand access. Fig. 4c presents an energy efficiency estimation by enhancing SOTA designs with the proposed hardware. On average, SATA can provide 1.34 \times energy efficiency and 1.3 \times throughput gain after integration. A³'s recursive search dominates runtime overhead and shows limited improvement.

V. CONCLUSION

We presented SATA, a sparsity-aware scheduling framework that enhances selective-token attention by improving operand locality, reducing redundant memory access, and sustaining high utilization. With tiling and zero-skip extensions, SATA scales efficiently to long sequences. Evaluations with silicon-validated simulations and selective-attention traces demonstrate throughput and energy efficiency improvements with low overhead. By bridging algorithmic sparsity with hardware utilization, SATA provides a locality-centric, proactive scheduling that complements existing accelerators and paves the way for more scalable and efficient Transformer deployment.

ACKNOWLEDGEMENTS

This work was supported in part by CoCoSys, one of seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

REFERENCES

- [1] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.
- [2] A. Dosovitskiy, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [3] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- [4] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International conference on machine learning*, pp. 10347–10357, PMLR, 2021.
- [5] Z. Wan, J. Qian, Y. Du, J. Jabbour, Y. Du, Y. Zhao, A. Raychowdhury, T. Krishna, and V. J. Reddi, "Generative ai in embodied systems: System-level analysis of performance, efficiency and scalability," in *2025 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 26–37, IEEE, 2025.
- [6] J. Devlin, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [7] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24824–24837, 2022.
- [8] A. Radford, "Improving language understanding by generative pre-training," 2018.
- [9] T. Xie, J. Zhao, Z. Wan, Z. Zhang, Y. Wang, R. Wang, R. Huang, and M. Li, "Realm: Reliable and efficient large language model inference with statistical algorithm-based fault tolerance," *arXiv preprint arXiv:2503.24053*, 2025.
- [10] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican, *et al.*, "Gemini: a family of highly capable multimodal models," *arXiv preprint arXiv:2312.11805*, 2023.
- [11] J. Lu, D. Batra, D. Parikh, and S. Lee, "Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks," *Advances in neural information processing systems*, vol. 32, 2019.
- [12] P. H. Seo, A. Nagrani, A. Arnab, and C. Schmid, "End-to-end generative pretraining for multimodal video captioning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17959–17968, 2022.
- [13] Z. Wan, Y. Du, M. Ibrahim, J. Qian, J. Jabbour, Y. Zhao, T. Krishna, A. Raychowdhury, and V. J. Reddi, "Reca: Integrated acceleration for real-time and efficient cooperative embodied autonomous agents," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pp. 982–997, 2025.
- [14] C. Wang, Z. Wan, H. Kang, E. Chen, Z. Xie, T. Krishna, V. J. Reddi, and Y. Du, "Slm-mux: Orchestrating small language models for reasoning," *arXiv preprint arXiv:2510.05077*, 2025.
- [15] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green ai," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.
- [16] S. Kim, C. Hooper, T. Wattanawong, M. Kang, R. Yan, H. Genc, G. Dinh, Q. Huang, K. Keutzer, M. W. Mahoney, *et al.*, "Full stack optimization of transformer inference: a survey," *arXiv preprint arXiv:2302.14017*, 2023.
- [17] M. Ibrahim, Z. Wan, H. Li, P. Panda, T. Krishna, P. Kanerva, Y. Chen, and A. Raychowdhury, "Special session: Neuro-symbolic architecture meets large language models: A memory-centric perspective," in *2024 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ ISSS)*, pp. 11–20, IEEE, 2024.
- [18] H. Fan, T. Chau, S. I. Venieris, R. Lee, A. Kouris, W. Luk, N. D. Lane, and M. S. Abdelfattah, "Adaptable butterfly accelerator for attention-based nns via hardware and algorithm co-design," in *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 599–615, IEEE, 2022.
- [19] T. J. Ham, Y. Lee, S. H. Seo, S. Kim, H. Choi, S. J. Jung, and J. W. Lee, "Elsa: Hardware-software co-design for efficient, lightweight self-attention mechanism in neural networks," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pp. 692–705, IEEE, 2021.
- [20] T. Dao, "Flashattention-2: Faster attention with better parallelism and work partitioning," *arXiv preprint arXiv:2307.08691*, 2023.
- [21] K. Roy, A. Kosta, T. Sharma, S. Negi, D. Sharma, U. Saxena, S. Roy, A. Raghunathan, Z. Wan, S. Spetalnick, *et al.*, "Breaking the memory wall: next-generation artificial intelligence hardware," *Frontiers in Science*, vol. 3, p. 1611658, 2025.
- [22] T. J. Ham, S. J. Jung, S. Kim, Y. H. Oh, Y. Park, Y. Song, J.-H. Park, S. Lee, K. Park, J. W. Lee, *et al.*, "A³: Accelerating attention mechanisms in neural networks with approximation," in *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 328–341, IEEE, 2020.
- [23] H. Wang, Z. Zhang, and S. Han, "Spatten: Efficient sparse attention architecture with cascade token and head pruning," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 97–110, IEEE, 2021.
- [24] Z. Zhou, J. Liu, Z. Gu, and G. Sun, "Energon: Toward efficient acceleration of transformers using dynamic sparse attention," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 1, pp. 136–149, 2022.
- [25] S. Tuli and N. K. Jha, "Acceltran: A sparsity-aware accelerator for dynamic inference with transformers," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 11, pp. 4038–4051, 2023.
- [26] G. Zhao, J. Lin, Z. Zhang, X. Ren, Q. Su, and X. Sun, "Explicit sparse transformer: Concentrated attention through explicit selection," *arXiv preprint arXiv:1912.11637*, 2019.
- [27] P. Wang, X. Wang, F. Wang, M. Lin, S. Chang, H. Li, and R. Jin, "Kvt: k-nn attention for boosting vision transformers," in *European conference on computer vision*, pp. 285–302, Springer, 2022.
- [28] X. Chen, H. Li, M. Li, and J. Pan, "Learning a sparse transformer network for effective image deraining," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5896–5905, 2023.
- [29] M. K. Singh, R. Yasarla, H. Cai, M. Lee, and F. Porikli, "Tosa: Token selective attention for efficient vision transformers," *arXiv preprint arXiv:2406.08816*, 2024.
- [30] Y. Xiao, Q. Yuan, K. Jiang, J. He, C.-W. Lin, and L. Zhang, "Ttst: A top-k token selective transformer for remote sensing image super-resolution," *IEEE Transactions on Image Processing*, 2024.
- [31] X. Fu, Q. Ren, H. Wu, F. Xiang, Q. Luo, J. Yue, Y. Chen, and F. Zhang, "P³ vit: A cim-based high-utilization architecture with dynamic pruning and two-way ping-pong macro for vision transformer," *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2023.
- [32] F. Tu, Z. Wu, Y. Wang, L. Liang, L. Liu, Y. Ding, L. Liu, S. Wei, Y. Xie, and S. Yin, "Trancim: Full-digital bitline-transpose cim-based sparse transformer accelerator with pipeline/parallel reconfigurable modes," *IEEE Journal of Solid-State Circuits*, pp. 1798–1809, 2022.
- [33] F. Tu, Z. Wu, Y. Wang, W. Wu, L. Liu, Y. Hu, S. Wei, and S. Yin, "Multcim: Digital computing-in-memory-based multimodal transformer accelerator with attention-token-bit hybrid sparsity," *IEEE Journal of Solid-State Circuits*, 2023.
- [34] X. Yang, B. Yan, H. Li, and Y. Chen, "Retransformer: Reram-based processing-in-memory architecture for transformer acceleration," in *Proceedings of the 39th International Conference on Computer-Aided Design*, pp. 1–9, 2020.
- [35] D. Bahdanau, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [36] G. Cheng, J. Han, and X. Lu, "Remote sensing image scene classification: Benchmark and state of the art," *Proceedings of the IEEE*, vol. 105, no. 10, pp. 1865–1883, 2017.
- [37] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, Ieee, 2009.
- [38] W. Yang, R. T. Tan, J. Feng, J. Liu, Z. Guo, and S. Yan, "Deep joint rain detection and removal from a single image," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [39] A. Lu, X. Peng, W. Li, H. Jiang, and S. Yu, "Neurosim simulator for compute-in-memory hardware accelerator: Validation and benchmark," *Frontiers in artificial intelligence*, vol. 4, p. 659060, 2021.
- [40] S. H. Xiao Chen Peng, "User manual of dnn-neurosim framework v2.0." https://github.com/neurosim/DNN_NeuroSim_V2.1, 2020.
- [41] R. Raj, S. Banerjee, N. Chandra, Z. Wan, J. Tong, A. Samajdar, and T. Krishna, "Scale-sim v3: A modular cycle-accurate systolic accelerator simulator for end-to-end system analysis," in *2025 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 186–200, IEEE, 2025.