



## Research

**Cite this article:** Glint T, Mittal B, Sharma S, Ronak AQ, Goud A, Kasture N, Momin Z, Krishna A, Mekie J. 2025 AxLaM: energy-efficient accelerator design for language models for edge computing. *Phil. Trans. R. Soc. A* **383**: 20230395.

<https://doi.org/10.1098/rsta.2023.0395>

Received: 4 June 2024

Accepted: 12 November 2024

One contribution of 11 to a theme issue 'Emerging technologies for future secure computing platforms'.

### Subject Areas:

artificial intelligence, electrical engineering, systems theory

### Keywords:

transformer accelerator, language model BERT, hardware accelerator

### Author for correspondence:

Joycee Mekie

e-mail: [joycee@iitgn.ac.in](mailto:joycee@iitgn.ac.in)

# AxLaM: energy-efficient accelerator design for language models for edge computing

Tom Glint<sup>1</sup>, Bhumika Mittal<sup>2</sup>, Santripta Sharma<sup>2</sup>, Abdul Qadir Ronak<sup>3</sup>, Abhinav Goud<sup>3</sup>, Neerja Kasture<sup>3</sup>, Zaqi Momin<sup>3</sup>, Aravind Krishna<sup>3</sup> and Joycee Mekie<sup>3</sup>

<sup>1</sup>Forschungszentrum Jülich, Jülich, Germany

<sup>2</sup>Ashoka University, Sonapat, Haryana, India

<sup>3</sup>Indian Institute of Technology Gandhinagar, Gandhinagar, Gujarat, India

JM, 0000-0001-9646-1941

Modern language models such as bidirectional encoder representations from transformers have revolutionized natural language processing (NLP) tasks but are computationally intensive, limiting their deployment on edge devices. This paper presents an energy-efficient accelerator design tailored for encoder-based language models, enabling their integration into mobile and edge computing environments. A data-flow-aware hardware accelerator design for language models inspired by Simba, makes use of approximate fixed-point POSIT-based multipliers and uses high bandwidth memory (HBM) in achieving significant improvements in computational efficiency, power consumption, area and latency compared to the hardware-realized scalable accelerator Simba. Compared to Simba, AxLaM achieves a ninefold energy reduction, 58% area reduction and 1.2 times improved latency, making it suitable for deployment in edge devices. The energy efficiency of AxLaM is 1.8 TOPS/W, 65% higher than FACT, which requires pre-processing of the language model before implementing it on the hardware.

This article is part of the theme issue 'Emerging technologies for future secure computing platforms'.

## 1. Introduction

The rapid advancements in natural language processing (NLP) have led to the development of powerful language models such as GPT and bidirectional encoder representations from transformers (BERT) [1,2], which excel in tasks such as text classification, question answering and language translation. However, the computational intensity and memory requirements of these models pose significant challenges for deployment on edge devices, which are constrained by limited power and computational resources [3].

Edge deployment of language models is crucial for applications requiring real-time processing and data privacy, such as mobile assistants, IoT devices and autonomous systems. Offloading computations to the cloud introduces latency and potential security risks owing to data transmission [4]. Therefore, there is a pressing need for specialized hardware accelerators that can efficiently run these models within the stringent power and area budgets of edge devices, compared to general-purpose computing environments.

In this paper, we present an energy-efficient accelerator design, for battery-powered edge devices, specifically optimized for encoder-based models such as BERT. Our main contributions are as follows:

(i) We, for the first time, show that a highly quantized and approximated variant of POSIT [5] can incur a negligible loss of accuracy compared to the Float-Point32 (FP32) encoding scheme for BERT. The resultant multiply accumulate (MAC) unit requires approximately 0.3 pJ for a MAC operation in 65 nm (approx. 69 times better than FP32). (ii) We present a fine-tuning method that allows the direct use of the pre-trained BERT to use the number system without structural changes to the neural network. To the best of our knowledge, this is the first time quantization without accuracy loss has been achieved without structural changes and retraining. (iii) Based on the hardware–software co-design approach, we introduce an optimized processing element (PE) architecture that leverages near-data processing on high bandwidth memory (HBM) [6] and the approximate fixed-point POSIT (AFPOS) number system to significantly reduce power consumption and area, by leveraging, for the first time, the three-dimensional spatial data reuse opportunity in BERT. (iv) We provide a comprehensive evaluation of our design against the hardware-realized accelerator Simba, demonstrating improvements in computational efficiency, power consumption, area and latency, while maintaining negligible accuracy loss on the CoLA dataset using BERT-large. Furthermore, we compare our work with eight recent accelerators for LLMs (that require structural modification to BERT) and demonstrate superior area and energy efficiency.

The rest of the paper is organized as follows: in §2 we review related work, in §3 we give details of our accelerator design approach, while in §4 we describe the methodology and evaluation infrastructure. In §5 we present the experimental results and analysis, then in §6 we discuss implications and limitations; §7 concludes the paper.

## 2. Related work

Hardware–software co-design has long been central to optimizing system performance by considering both hardware and software components. Early research, such as [3] and [7], emphasized automatic parallelization and self-programming systems to reduce memory operations and enhance data communication. In machine learning accelerators, especially for deep learning, several architectures have emerged. Simba [8] introduced a scalable multi-chip framework optimized for convolutional neural networks (CNNs), but it faces challenges with transformer models owing to differing computational patterns.

Recent accelerators specifically target transformers. For instance, OPTIMUS [9] focuses on enhancing matrix multiplication efficiency, while  $A^3$  [10] accelerates the attention mechanism through the use of approximation techniques. NVIDIA's Transformer Engine [11] employs specialized tensor cores for faster inference, and SwiftTron [12] applies 8-bit

quantization for resource-constrained environments. Further advancements include co-optimization and processing-in-memory (PIM) strategies. The H3D-Transformer [13] combines PIM with edge-specific optimizations, while ReTransformer [14] utilizes a ReRAM-based PIM architecture. TransPIM [15] focuses on memory-based acceleration through hardware–software co-design. More details on these comparisons are provided in §5, specifically in §5c. *Current designs often require modifications to neural networks or lack comprehensive system-level implementations, particularly in managing off-chip memory access and supporting various integer formats. Moreover, these accelerators typically necessitate significant changes to the model structure to fully support transformers.*

Our work introduces an accelerator leveraging HBM3 and the AFPOS number system, optimized for encoder-based models such as BERT. This design offers a system-level solution for edge deployment, improving performance and efficiency without altering the model structure. In addition, number system-aware fine-tuning of the BERT model ensures accuracy comparable to FP32. For hardware comparison, Simba is used owing to its flexibility and broad acceptance, as recent works lack silicon implementation. Nonetheless, results are evaluated against recent accelerators, as discussed in §5, particularly in §5c.

### 3. Accelerator design approach

#### (a) Background—BERT: use cases and internal structure

BERT [1] is a pre-trained language model introduced by Google in 2018. Unlike traditional CNN models, BERT uses bidirectional training, leveraging both preceding and succeeding contexts to predict words, which results in a deeper understanding of language. This approach has led to state-of-the-art performance across various NLP tasks. BERT excels in tasks such as text classification, named entity recognition, question answering, sentiment analysis and language translation [16]. Its pre-trained nature allows for fine-tuning, making it versatile and well-suited for edge applications.

BERT is based on the Transformer architecture, consisting of an input embedding layer, multiple Transformer encoder blocks and an output layer. Each encoder block includes a multi-head self-attention mechanism and a feed-forward neural network. Tokens are embedded with positional encodings and processed through Transformer blocks. The BERT-large model has specific parameters: a model dimensionality ( $d_{\text{model}}$ ) of 1024, 24 Transformer blocks ( $N$ ), each with 16 attention heads ( $H$ ). The key/query ( $d_k$ ) and value ( $d_v$ ) vectors have dimensions  $d_{\text{model}}/H = 64$ . Each attention head computes a  $1024 \times 1024$  matrix, with concatenated outputs passed through a linear transformation to produce the final layer output:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \text{head}_2, \dots)W_O,$$

where  $W_O$  is the weight matrix that maps the concatenated output to the embedding size,  $d_{\text{model}} = 1024$ . After multi-head attention, the output undergoes a residual connection and layer normalization and is processed through a feed-forward neural network, which first increases and then reduces dimensionality. The matrix operations are listed in table 1. These operations enhance computational efficiency and consistency. BERT's effect on NLP tasks has made it a leading model, with scalable configurations such as BERT-base and BERT-large suitable for various environments. BERT's performance is recognized in *MLPerf* [16], making it the baseline workload for our design.

#### (b) Background—spatial and temporal reuse in hardware accelerator

While efforts exist to accelerate BERT [17,18], no accelerators have been specifically optimized for edge scenarios, as noted in [2,3]. In edge computing, the Simba architecture is a leading solution for deep neural network inference [8,19], thanks to its advanced vector multipliers

**Table 1.** Summary of unique matrix operations in BERT's encoder.

operation	matrix $L$ dimensions	matrix $R$ dimensions	result dimensions
$Q/K/V$	$1024 \times 1024$	$1024 \times 64$	$1024 \times 64$
attention	$1024 \times 64$	$64 \times 1024$	$1024 \times 1024$
multi-head	$1024 \times 1024$	$1024 \times 1024$	$1024 \times 1024$
bottleneck expand	$1024 \times 1024$	$1024 \times 4096$	$1024 \times 4096$
bottleneck contract	$1024 \times 4096$	$4096 \times 1024$	$1024 \times 1024$

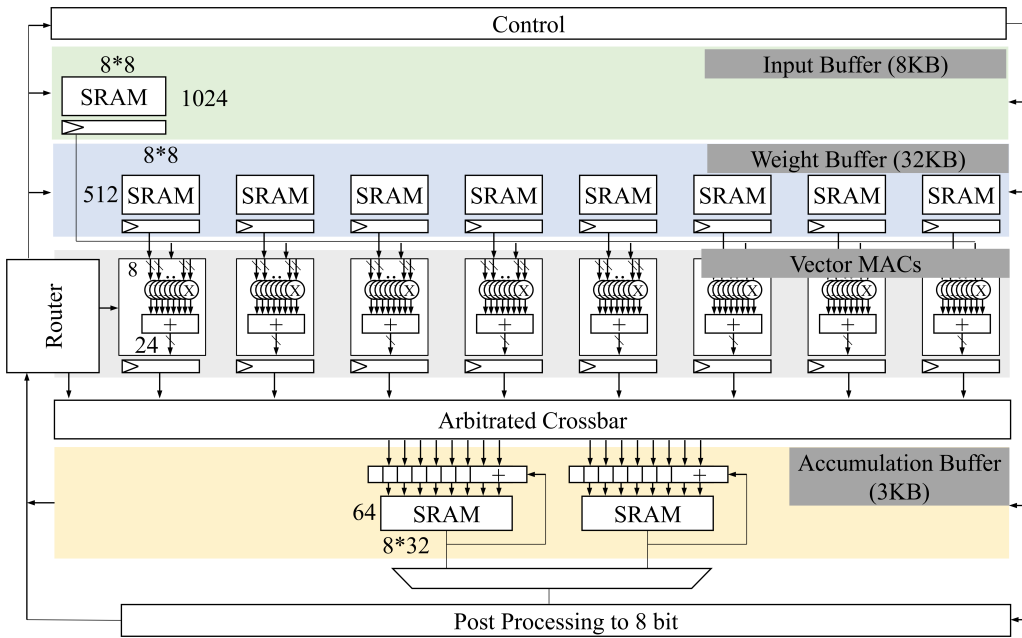
and accumulators that share input activations while outputting to multiple channels. Notably, Simba features a data-agnostic design, making it robust against data distribution variations.

At the core of Simba, illustrated in figure 1, is the PE, responsible for executing convolutional layers, fully connected layers and post-processing operations such as bias addition, ReLU and pooling. Each PE contains eight lanes, each using a distinct weight tensor to generate elements for a single output channel ( $K$ ). An 8-bit precision vector MAC unit within each lane multiplies eight input elements from different channels ( $C$ ) with eight weight elements, summing them to produce a single output. This approach is highly efficient, given typical channel counts (64–1024). The architecture is supported by local input activation, output activation and accumulation SRAMs, which buffer the datapath. Energy efficiency is optimized by minimizing SRAM access: input activation SRAM is read every cycle, and its elements are forwarded across lanes. The weight SRAM, wider than the input activation SRAM, supplies distinct vectors to each lane and reuses weights across multiple inputs, reusing them  $P \times Q$  times (where  $P$  and  $Q$  are output dimensions). Accumulation SRAM stores intermediate sums, conserving energy by accumulating the eight-wide vector from  $C$  channels. The output size, often exceeding the accumulation buffer's capacity, requires temporal tiling to generate portions of output activations sequentially. This buffer also supports cross-PE reductions when the weight kernel spans multiple PEs. Post-processing tasks such as ReLU, bias addition and pooling are performed after accumulation. The buffer is dual-banked, allowing simultaneous access by MACs, routers and post-processing units, with an arbitration crossbar to resolve bank conflicts.


### (c) Motivation for a new architecture design

The growing complexity and computational demands of NLP, particularly models such as BERT, have exposed the limitations of existing hardware accelerators for machine learning (ML). Accelerators for BERT, which are primarily based on FPGA platforms [18] or the more recent ASIC designs [2], often require significant modifications to the neural network architecture to achieve operational efficiency. Additionally, such modifications, including extreme quantization or mixed-signal processing techniques, frequently result in a notable loss of accuracy [2,13,20]. While the Simba architecture, owing to silicon realization and data-agnostic operation, has been recognized as a leading solution for deep neural network inference, its efficacy diminishes when confronted with the distinct computational demands of BERT's matrix operations. The various matrix multiplications, particularly those involving the QKV (Query, Key, Value) operations, impose a substantial power burden, primarily owing to the significant DRAM accesses required, as illustrated in figure 2. These multiplications, as listed in table 1, highlight the computational challenges that Simba faces when deployed for models such as BERT.

The QKV operations, though not the largest in magnitude, play a critical role in BERT's performance owing to relatively frequent computation (16 times each per encoder block). The power consumption associated with these operations is exacerbated by the limited data reuse inherent in their matrix multiplications. In energy-efficient computing, small buffer sizes are



**Figure 1.** Simba's PE architecture [8].



**Figure 2.** Simba's (with FP16 MAC for accuracy preservation) system power for processing matrices in BERT's encoder.

often preferred to optimize the power consumption of such operations. However, as the buffer sizes increase, the energy required for data access also rises, leading to further inefficiencies. Conversely, operations such as Bottleneck Contraction (BtlCont), Bottleneck Expansion (BtlExp) and Multi-Head Attention (MHead) exhibit a more pronounced data reuse pattern. Each output word in these operations depends on approximately 1024 input words, allowing for a significant reduction in access counts as buffer sizes grow. This characteristic makes large buffer sizes essential to fully exploit data reuse, thereby improving power efficiency in these layers.

Despite these insights, the Simba architecture's current design is not well-suited to meet the diverse computational requirements of BERT's matrix operations. Simba's average power consumption, 9 and 19 W, respectively, for INT8 and FP16 pipelines, is far beyond the acceptable range for edge and portable applications, which typically impose stringent power constraints, often below 5 W [21]. Therefore, while Simba excels in certain scenarios, its inability to efficiently handle the unique demands of BERT's matrix multiplications underscores the need for a new architecture that can better address these challenges.

## (d) Design consideration: challenges to be addressed


Efficient hardware accelerators for neural networks such as BERT must address several key challenges: high DRAM access count, on-chip buffer energy consumption and on-chip compute energy efficiency.

- (i) Reducing high DRAM access count and energy consumption: high DRAM access significantly effects overall energy consumption. To mitigate this, increasing the on-chip buffer size can reduce the frequency of DRAM accesses by enhancing data reuse and minimizing energy-intensive memory transactions. This approach leverages larger buffers to store more data locally, thereby reducing the need for frequent access to external memory.
- (ii) Reducing energy at the on-chip buffer level: while larger on-chip buffers can reduce DRAM access, they can also lead to increased energy consumption within the chip. This challenge can be managed by optimizing memory hierarchies with multi-level buffering strategies and exploiting opportunities for temporal and data sharing within the application. Effective buffer management ensures that the energy benefits of reduced DRAM access are not offset by increased internal energy costs.
- (iii) The energy required for on-chip computations, particularly MAC operations, is a significant consideration for BERT's intensive computations. Reducing this energy footprint can be achieved by employing low-power arithmetic units with reduced precision and by implementing operand-sharing techniques to maximize the reuse of input data. These strategies help lower the energy consumption associated with each computation, making the overall processing more efficient.

## (e) Design consideration: on-chip buffer size and memory accesses

The unique matrix multiplications within BERT's encoder exhibit distinct DRAM access patterns that significantly effect the power consumption associated with different buffer sizes. The primary objective is to select a buffer size that minimizes the overall system energy consumption (considering on-chip data reuse) during data accesses and computational processes. As shown in [figure 3a](#), the DRAM access count varies with increasing total buffer size, offering valuable insights into data reuse and how frequently each matrix multiplication operation requires access to main memory. To complement this, [figure 3b](#) illustrates the energy per byte access data for different buffer sizes, highlighting the energy implications of choosing specific buffer configurations. Balancing the frequency of DRAM accesses against the energy overhead of larger buffers is critical in determining the most energy-efficient buffer size. To achieve optimal energy efficiency, both the frequency of DRAM accesses for each matrix multiplication and the energy cost associated with each access—given a specific buffer size—must be considered. Additionally, the repetitive nature of BERT's matrix operations within the encoder should be taken into account, as certain operations recur more frequently, potentially skewing the overall energy profile. For instance, while QKV operations may require fewer DRAM accesses in some buffer configurations, their central role in the encoder and frequent occurrence contribute significantly to the total energy consumption.

Comparing the DRAM access count data with buffer access energy, it becomes evident that smaller buffer sizes generally exhibit lower energy costs per access. However, as buffer sizes increase, while DRAM access counts typically decrease, the energy cost per access rises. Considering all available metrics, an intermediate buffer size appears to strike the best balance between reducing DRAM accesses and maintaining manageable energy costs per access. Analysis suggests that buffer sizes in the range of 64–128 KB may be appropriate, as indicated by the trends in the provided data. However, further refinement of this range is necessary, particularly by considering the specific frequency of each matrix operation in BERT's encoder.



**Figure 3.** (a) DRAM access count versus on-chip buffer size while processing BERT. (b) Access energy of buffers against size, compute and HBM.

## (f) Opportunity: HBM logic layer for bandwidth and low access energy

Traditional DDR DRAM (such as DDR4 DRAM) is insufficient for handling the high frequency memory accesses required by BERT's encoders, where DRAM access energy is a major concern [22]. HBM3 [6] offers a tenfold reduction in access energy per bit (4.2 versus 46 pJ), making it a superior choice for BERT's frequent matrix multiplications, significantly enhancing accelerator efficiency. HBM3's design is well-suited for BERT-like models, allowing direct integration of accelerator designs within its logic layer (up to 30 mm<sup>2</sup> and 8.5 W power [23]), as supported by recent manufacturer initiatives [24]. This integration minimizes data movement, reducing latencies and further cutting energy consumption, crucial for the high data reuse in BERT's encoders. HBM3's structure stacks DRAM dies atop a logic layer, connected via through-silicon vias, achieving a compact memory footprint. It vastly outperforms DDR4 in bandwidth, offering approximately 512 GB/s per stack compared to DDR4's 25.6 GB/s, making HBM3 ideal for high-throughput, rapid-access applications [6]. However, this presents three main challenges: ensuring balanced access to all HBM memory channels, managing design area constraints and addressing thermal and power limitations.

## (g) Addressing the high MAC energy

In accelerators for complex language models such as BERT, MAC operations are the second highest contributors to power consumption after DRAM access, as shown in figure 2. Traditional 8-bit integer operations, while efficient, do not meet the precision demands of BERT, necessitating a solution that combines floating-point precision with reduced power overheads (figure 4). Recent advancements in multiplier design, such as the POSIT number system introduced by Gustafson & Yonemoto [5], offer a promising alternative. POSITs, defined by parameters ( $N, es$ ), provide dynamic flexibility, allowing efficient space usage by encoding exponent and fraction parts only when necessary. However, the increased area and power requirements of POSITs present challenges. To mitigate these, the fixed-POSIT representation was developed [25], which standardizes the regime length ( $rf$ ) alongside  $N$  and  $es$ . The AFPOS multiplier further refines this approach, optimizing for area, energy and latency while requiring 25% less bit storage than fixed-POSITs [22]. Our experimentation has shown that fixing the  $2^{\beta}$  term for AFPOS (with 1-bit sign, 4-bit exponent and 3-bit mantissa) to  $2^{-7}$  results in negligible loss of accuracy compared to FP32, as shown in figure 5. For the rest of the paper, AFPOS refers to this ( $N = 10, es = 4, \beta = -7$ ) configuration. Table 2 lists the latency, area and energy of these multipliers. We have highlighted the designs with the highest quantization yet having similar accuracy and Matthews correlation coefficient (MCC) score as FP32. Compared to a quantized 16-bit floating-point multiplier, AFPOS reduces area by a factor of 4.04, energy by a factor of 8.9

and latency by a factor of 2.1. We use a 16-wide vector MAC, with area, latency and power as listed in table 3, for efficient spatial reduction.

## (h) PE design details and innovations

BERT's computational architecture, particularly its vector multipliers, poses unique challenges distinct from frameworks such as Simba. In Simba, shared input activations among a set of vector multipliers, along with exclusive weights, allow weights to be read once and reused multiple times. This resource optimization, however, is not inherent in BERT. When dealing with matrices  $L$  and  $R$  (as seen in table 1), if the elements of  $R$  are temporarily stationary, those of  $L$  require constant cycling and vice versa. This dynamic underscores the need to efficiently amortize read energy costs, leading to a redesigned vector MAC and PE, as shown in figure 6. Instead of employing 16 separate PEs, we propose a unified PE that optimizes shared resource utilization while maintaining an acceptable fan-out. For each read of an element in  $L$ , the element contributes to the partial product of eight distinct columns, while efficient reads from  $R$  enable participation in eight different rows of the output matrix. A simplified representation of this concept is illustrated in figure 7, where words read together are annotated with alphanumeric labels and sources are colour-coded equally compared to figure 6. Furthermore, within a column or row, 16 elements are multiplied together and accumulated using the vector MAC and written to the accumulation buffer (A.SRAM). We choose this configuration, as the baseline Simba instance, in efficiency mode, offers 1 TOPS throughput (which offers sub-1 s inference of BERT-large). Our goal is to achieve a similar throughput with our vector MAC configuration, with the parallel adder tree. The operational strategy dictates that either the  $L$  or  $R$  matrix elements must be refreshed upon completing each operand set. To balance energy efficiency and performance, our design incorporates an 8 KB buffer for all  $L$  buffers and an equivalent capacity for the  $R$  buffers. The A.SRAM is designed to cache partial products, enabling efficient summation.

### (i) Design parameter justification

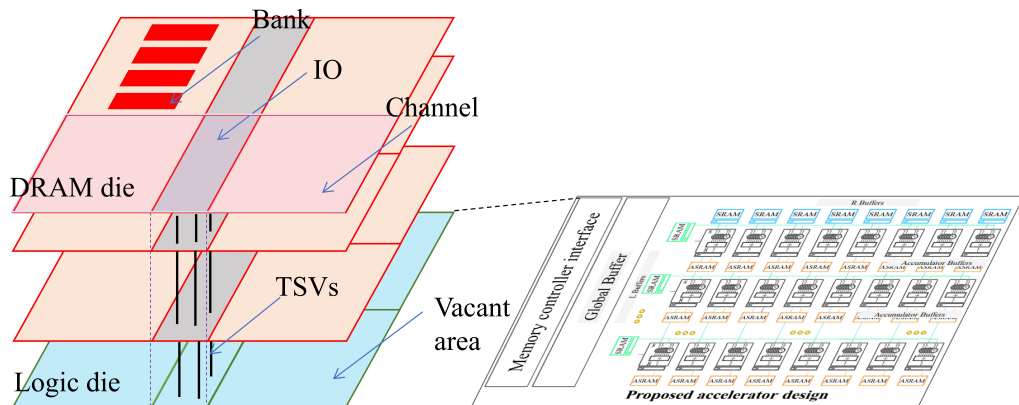
The unified PE is directly interfaced with all 16 channels of HBM through memory controllers, occupying only 2.5 mm<sup>2</sup> of the available 30 mm<sup>2</sup> space, leaving ample room for additional designs, as shown in figure 4. The vector MAC operates at 500 MHz (as listed in table 3) to optimize energy efficiency. The choice of a 16-wide vector MAC is based on achieving the desired throughput of 1 TOPS, equivalent to Simba's efficiency mode. Further widening of the vector MAC would introduce significant combinational delays, while increasing the operating frequency would substantially raise area and energy consumption. This throughput also meets the sub-1 s response time expected for edge devices. Simulations revealed that increasing the total number of multipliers led to diminishing returns owing to memory bandwidth limitations. The 8 KB buffer size was selected to strike a balance between on-logic die storage and energy consumption. While larger buffers reduce DRAM accesses, they also increase access energy per byte, as shown in figure 3b. Our analysis determined that 8 KB of local buffer per row and column offers an optimal balance for handling BERT's workload efficiently.

## 4. Methodology


### (a) Evaluation infrastructure

At the software level, we model BERT using a PyTorch-based GPU accelerated framework, modified to support different number representations, including AFPOS. In the framework, the flow of data is altered such that the data are encoded into the specified encoding scheme (such as BFloat, POSIT or AFPOS) before multiplication. This allows the framework to use the IEEE





**Figure 4.** HBM3 structure showing vacant space for accelerator placement.



**Figure 5.** Model accuracy and MCC score of BERT for various encoding schemes, fine-tuned on the CoLA dataset. Approximate fixed POSIT value =  $(-1)^{\text{sign}} \times 2^{\beta} \times 2^{\text{exp}} \times (1.\text{mantissa})$

754 Float-based pre-trained BERT model without alteration. After fine-tuning the model with the CoLA dataset, we measure the accuracy and MCC scores against the full-precision model (result shown in figure 5). We use this software framework for reproducibility and for exact computation and score calculation.

For hardware modelling, we use Cadence Genus to synthesize the MAC units (of all encoding schemes) at a 65 nm node, ensuring timing constraints are met at 500 MHz frequency. SRAM buffers are modelled using CACTI [26]. System-level modelling and optimal data search (for least latency followed by least energy) are performed with Timeloop, Accelergy and Alladin framework [7,27], defining both Simba and our proposed architecture based on synthesized values. For Simba, the external DRAM memory used is DDR4 DRAM, and the hardware model is identical to the hardware-realized version of Simba [8]. Our proposed work is modelled as a design instantiated on the logic layer of the HBM3 memory. We use this hardware model framework as it has been validated against real hardware and allows for flexible extensions.

## (b) Workload description

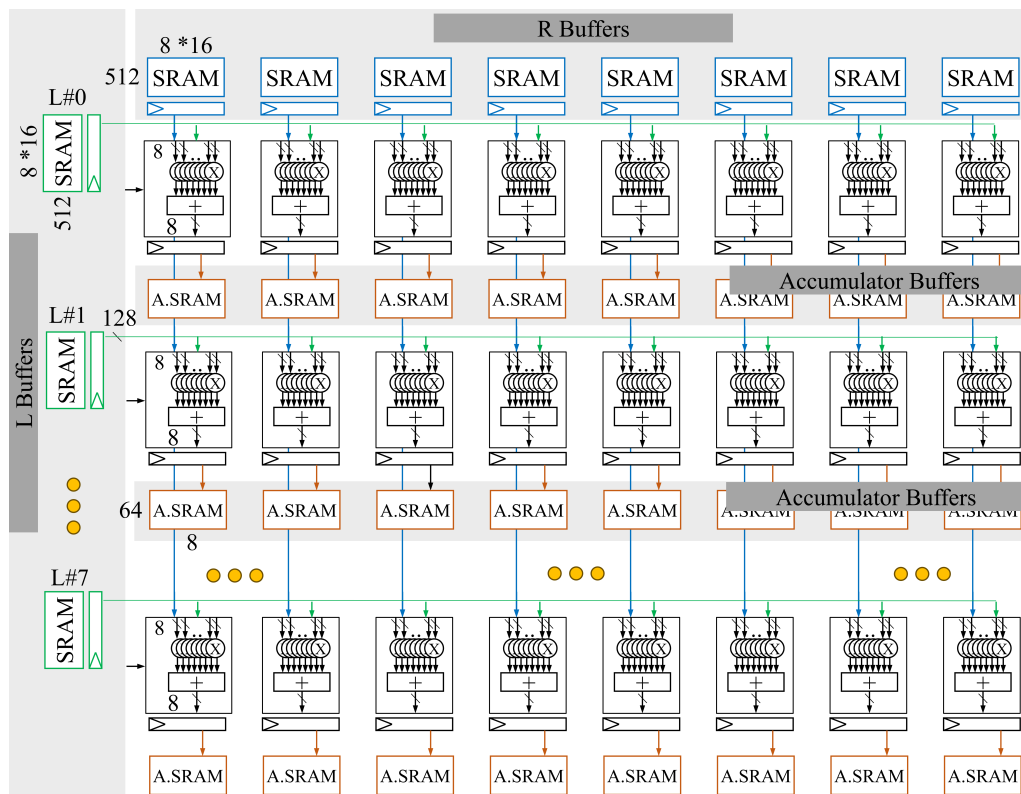
We focus on the unique matrix operations within BERT's encoder, as listed in table 1. These operations include the QKV computations, attention mechanisms, multi-head concatenations and feed-forward networks. The network level performance of any variation of BERT-large can be extrapolated from these matrices.

**Table 2.** Minimum latency of design, and power and area of multipliers against different data formats.


type	bits	config.	latency (ps)	area ( $\mu\text{m}^2$ )	energy ( $\mu\text{J}$ )	type	bits	config.	latency (ps)	area ( $\mu\text{m}^2$ )	energy ( $\mu\text{J}$ )
FP32	32		1299	13830.84	22.50	FPOS	14	14,4,2	558	1003.68	1.68
FP16	16		979	3095.28	4.55	FPOS	14	14,2,3	656	991.08	1.12
BF16	16		832	2148.12	2.75	FPOS	13	13,4,2	528	1047.24	1.82
POSIT	16	16,3	1906	8634.60	13.45	FPOS	13	13,2,3	656	982.08	1.62
POSIT	16	16,2	1925	9327.60	14.84	FPOS	11	11,4,2	533	914.04	2.01
POSIT	8	8,3	1808	1062.72	1.58	FPOS	11	11,2,3	680	904.32	1.90
POSIT	8	8,2	1224	2508.48	5.08	FPOS	10	10,6,2	1004	480.24	0.51
POSIT	6	6,3	775	965.88	2.90	FPOS	10	10,4,2	999	409.32	0.48
FPOS	16	16,6,2	1225	597.96	0.61	FPOS	10	10,2,6	827	1578.24	1.11
FPOS	16	16,4,2	1114	510.84	0.60	FPOS	10	10,2,3	636	919.08	1.50
FPOS	16	16,2,6	826	1723.68	1.76	FPOS	9	9,4,2	532	898.92	1.73
FPOS	16	16,2,3	814	851.40	1.40	FPOS	9	9,2,3	640	774.36	1.55
FPOS	12	12,6,2	797	582.48	0.89	FPOS	8	8,4,2	550	524.88	1.18
FPOS	12	12,4,2	797	502.92	0.92	FPOS	8	8,2,3	648	790.20	0.74
FPOS	12	12,2,6	799	1459.80	1.26	FPOS	7	7,2,3	629	709.92	1.63
FPOS	12	12,2,3	796	681.84	0.86	APOS	8	10,4	465	766.08	0.51
FPOS	15	15,4,2	797	576.00	1.01	APOS	6	10,4	435	628.56	0.26
FPOS	15	15,2,3	675	1112.04	1.63						

**Table 3.** Synthesis result: area, energy and latency values of AFPOS vector MAC (16 multipliers and adder tree).

unit	constraint	area ( $\mu\text{m}^2$ )	timing (ps)	power (with stimulus) (W)			total
				switching	internal	dynamic	
AFPOS (6,2)	max freq	$2.38 \times 10^3$	$5.75 \times 10^2$	$2.11 \times 10^5$	$1.20 \times 10^4$	$1.41 \times 10^4$	$1.41 \times 10^4$
AFPOS (8,4)	max freq	$8.53 \times 10^3$	$9.56 \times 10^2$	$1.09 \times 10^3$	$1.45 \times 10^3$	$2.54 \times 10^3$	$2.54 \times 10^3$
AFPOS (10,4)	max freq	$1.55 \times 10^4$	$1.96 \times 10^3$	$1.22 \times 10^3$	$1.43 \times 10^3$	$2.65 \times 10^3$	$2.65 \times 10^3$



**Figure 6.** Unified PE design of the proposed accelerator.




**Figure 7.** Simplified example showing operations that can be performed in a single cycle in parallel.

## 5. Results

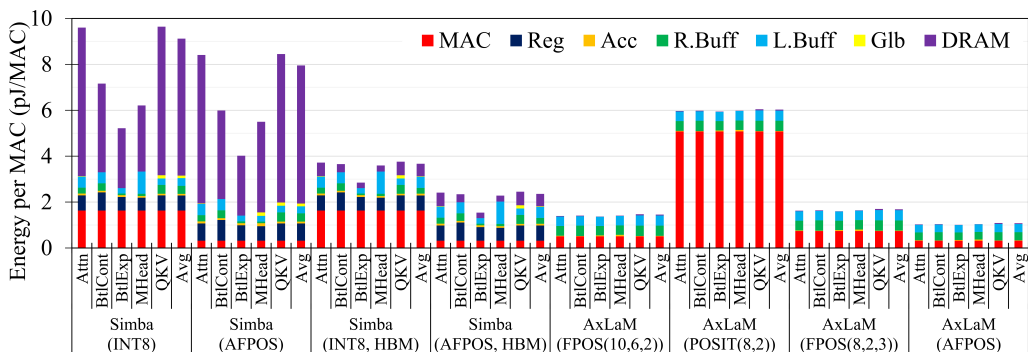
### (a) Latency and energy

Figure 8 compares the latency during the processing of each unique matrix by Simba and the proposed design. Although both designs feature an identical RAW throughput of 1 TOPS, Simba's reduced data reuse, particularly in the QKV and attention matrices, limits its ability to fully utilize all MAC units owing to bandwidth constraints. This inefficiency results in a 1.2 times speedup for the proposed design. Each encoder is processed within 27 ms, enabling the complete BERT-large model to be processed in 0.6 s, well within the acceptable latency threshold for handheld/edge devices performing natural language tasks, even when processing up to 1000 tokens.


Figure 9 shows the amortized energy cost of a MAC operation at the system level, comparing Simba, a variant of Simba (Simba with AFPOS multipliers, Simba implemented in HBM3's logic layer and Simba implemented with AFPOS multiplier in HBM3's logic layer) and the proposed designs implemented with HBM3 and the candidate multipliers. The average of each implementation in figure 9 represent weighted average operations corresponding to an encoder in BERT-large. The AFPOS-adapted version provides insight into the effect of changing



**Figure 8.** Performance comparison.



**Figure 9.** Energy comparison.



**Figure 10.** Simba versus proposed design area comparison.

the number system alone. The proposed design (AxLaM (AFPOS)) demonstrates the lowest energy per MAC operation with a ninefold reduction in energy consumption, primarily owing to decreased DRAM access and multiplication energy compared to Simba. By only altering the multiplier in Simba, the energy difference narrows to sevenfold for the proposed design. Additionally, the internal buffer energy is halved in the proposed design, which is attributed to the improved PE organization.

### (b) Area and power constraints comparison: Simba versus proposed design

Integrating the proposed design within the available area of the HBM3 memory's logic die offers substantial benefits in bandwidth and energy efficiency. Figure 10 illustrates the area distribution for the various components of the design. Compared to the Simba design, the proposed design not only adheres to the 36 mm<sup>2</sup> area constraint of HBM3 but also significantly reduces it to just 2.5 mm<sup>2</sup>. This efficiency is primarily due to the streamlined single PE design of the proposed system, as opposed to Simba's more expansive 16 PE architecture, resulting in an impressive 58% reduction in area.

In terms of power consumption, the proposed design shows marked improvement, consuming only 1.103 W, which is a ninefold reduction compared to Simba's 9.34 W (INT8).

**Table 4.** Comparison with recent accelerators.

accelerator	original implementation detail	target models	area (mm <sup>2</sup> )	throughput (TOPS)	normalized E. efficiency (TOPS/W)	remarks
OPTIMUS [9]	ASIC simulation, 28 nm synthesized circuits	transformers	5.2	0.5	0.127	INT 16
A3 [10]	ASIC simulation, 40 nm synthesized circuits	BERT base	2.08	N/A	N/A	INT 8
Transformer engine [11]	hardware implementation inside 4 nm H100 GPU	ANY	N/A	1978	0.04	FP 8
Swiftron [12]	65 nm ASIC synthesized and simulated	RoBERTa	273	8	0.236	INT 8
FACT [20]	28 nm ASIC synthesized and simulated	BERT	6.05	0.928	1.1*	INT 4 + 8 (HW–SW co-design)
ReTransformer [14]	28 nm RRAM-based CIM simulated	custom	N/A	0.081	0.2	MIXED SIGNAL
TransPIM [15]	22 nm DRAM-based CIM simulated	RoBERTa	53.15	0.734	N/A	BIT-SERIAL
H3D [13]	7 nm digital and 22 nm analogue CIM simulated	BERT/GPT	47.3	1.6	1.07	MIXED SIGNAL
Simba [8]	16 nm hardware implementation	vector × matrix	6	0.3 to 4	0.039–0.31	INT8 (silicon realized)
<b>AxLaM (OURS)</b>	<b>ASIC simulation, 65 nm synthesized circuits</b>	<b>BERT</b>	<b>2.5</b>	<b>1</b>	<b>1.85</b>	<b>AFPOS</b>

This significant decrease ensures that the proposed design comfortably fits within HBM3's 8.5 W power limit [6]. Power consumption is calculated based on the energy and time required to process one full encoder of BERT.

### (c) Comparison with other recent accelerators

Table 4 lists a comparison between our proposed design, AxLaM and recent accelerators such as OPTIMUS, A<sup>3</sup>, Transformer Engine, SwiftTron and others. Energy efficiency has been normalized to a 65 nm node based on voltage and capacitance scaling models. AxLaM distinguishes itself through the use of the AFPOS system, achieving a throughput of 1 TOPS with a normalized energy efficiency of 1.3 TOPS/W. This positions AxLaM as one of the most energy-efficient designs, particularly for BERT, all within a compact 2.5 mm<sup>2</sup> area on a 65 nm process node.

While other accelerators, such as the Transformer Engine, offer higher raw throughput, they fall short in terms of energy efficiency. AxLaM also demonstrates superior area efficiency, especially when compared to designs such as Swiftron and H3D, which require significantly more space. AxLaM effectively balances throughput, energy efficiency and area, making it a highly competitive solution for edge devices focused on NLP tasks.

In terms of energy efficiency, H3D [13] and FACT [20] are the closest competitors. However, FACT requires modifications to the language model. The actual benefit comes from algorithm changes applied to BERT before porting it on the hardware. On the other hand, H3D, being based on analogue computing, suffers from high variability and is not suitable for mass

manufacturability. This further emphasizes AxLaM's robustness and suitability for practical deployment.

## 6. Discussion

Our results demonstrate that the proposed accelerator design significantly improves energy efficiency, area utilization and latency for running language models such as BERT on edge devices. By integrating with HBM3 and utilizing the AFPOS number system, we address the key challenges of power consumption and memory access bottlenecks.

The accuracy of our AFPOS BERT-large model remains comparable to the full-precision model, ensuring that performance is not sacrificed for efficiency. The design scales well with input token sizes, making it suitable for a range of NLP applications.

However, our study is limited by the availability of detailed data for some recent accelerators, and actual hardware implementation is required to fully validate the simulation results. Future work includes implementing the design on an FPGA or ASIC platform and exploring further optimizations. The exploration and analysis of encoding schemes, including all variations of POSITs, though performed, have been left out for brevity.

## 7. Conclusion

We have presented an energy-efficient accelerator design tailored for encoder-based language models such as BERT, suitable for deployment on mobile and edge devices. By leveraging near-data processing, an optimized PE architecture and the AFPOS number system, our design achieves significant improvements over the hardware-realized state-of-the-art accelerator Simba. We also show the energy and area improvements over other LLM accelerators that require extensive DNN structural changes.

Our work contributes to enabling real-time, privacy-preserving NLP applications on edge devices, opening avenues for further research in efficient hardware designs for complex language models.

**Data accessibility.** Data are available on Zenodo [28].

**Declaration of AI use.** We have not used AI-assisted technologies in creating this article.

**Authors' contributions.** T.G.: investigation, methodology, writing—original draft, writing—review and editing; B.M.: investigation; S.S.: methodology; A.Q.R.: methodology, validation; A.G.: software; N.K.: software; Z.M.: validation; A.K.: investigation; J.M.: conceptualization, project administration, resources, software, supervision, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

**Conflict of interests.** We declare we have no competing interests.

**Funding.** This work was supported by the Science and Engineering Research Board (SERB), Government of India, under the SERB-SUPRA grant SPR/2020/000450, the Ministry of Electronics and Information Technology (MEITY) under the Chips to Startup (C2S) programme and the Semiconductor Research Corporation (SRC) through contract 2020-IR-2980. Additional support was provided by the Federal Ministry of Education and Research (BMBF, Germany) under the NEUROTEC II project (project no. 16ME0398K) and the NeuroSys project (grant no. 03ZU1106AB).

## References

1. Devlin J, Chang MW, Lee K, Toutanova K. 2018 Bert: Pre-Training of Deep Bidirectional Transformers for Language Understanding. *arXiv Preprint arXiv*

2. Koilia N, Kachris C. 2024 Hardware Acceleration of Llms: A Comprehensive Survey and Comparison *arXiv Preprint arXiv*.
3. Akkad G, Mansour A, Inaty E. 2024 Embedded Deep Learning Accelerators: A Survey on Recent Advances. *IEEE Trans. Artif. Intell.* 1–19. (doi:10.1109/taai.2023.3311776)
4. Alwarafy A, Al-Thelaya KA, Abdallah M, Schneider J, Hamdi M. 2021 A Survey on Security and Privacy Issues in Edge-Computing-Assisted Internet of Things. *IEEE Internet Things J.* 8, 4004–4022. (doi:10.1109/jiot.2020.3015432)
5. Gustafson J, Yonemoto I. 2017 Beating Floating Point at its Own Game: Posit Arithmetic. *JSFI* 4, 71–86. (doi:10.14529/jsfi170206)
6. Park MJ *et al.* 2023 A 192-Gb 12-High 896-GB/s HBM3 DRAM With a TSV Auto-Calibration Scheme and Machine-Learning-Based Layout Optimization. *IEEE J. Solid State Circuits* 58, 256–269. (doi:10.1109/jssc.2022.3193354)
7. Parashar A *et al.* 2019 Timeloop: A Systematic Approach to DNN Accelerator Evaluation. In 2019 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Madison, WI, USA, pp. 304–315. (doi:10.1109/ISPASS.2019.00042)
8. Zimmer B *et al.* 2020 A 0.32–128 TOPS, Scalable Multi-Chip-Module-Based Deep Neural Network Inference Accelerator With Ground-Referenced Signaling in 16 nm. *IEEE J. Solid State Circuits* 55, 920–932. (doi:10.1109/jssc.2019.2960488)
9. Park J, Yoon H, Ahn D, Choi J, Kim JJ. 2020 OPTIMUS: Optimized Matrix Multiplication Structure for Transformer Neural Network Accelerator. In *Proceedings of Machine Learning and Systems*, vol. 2, pp. 363–378,
10. Ham TJ *et al.* 2020 A<sup>3</sup>: Accelerating Attention Mechanisms in Neural Networks with Approximation. In 2020 IEEE International Symposium on High Performance Computer Architecture (HPCA), San Diego, CA, USA, pp. 328–341. (doi:10.1109/HPCA47549.2020.00035). <https://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=9044284>.
11. Corporation N. 2023 NVIDIA H100 tensor core GPU architecture overview. See <https://resources.nvidia.com/en-us-tensor-core> (accessed 6 April 2024).
12. Marchisio A, Dura D, Capra M, Martina M, Masera G, Shafique M. 2023 SwiftTron: An Efficient Hardware Accelerator for Quantized Transformers. In *International Joint Conference on Neural Networks (IJCNN)*, Gold Coast, Australia, pp. 1–9. (doi:10.1109/IJCNN54540.2023.10191521)
13. Luo Y, Yu S. 2024 H3D-Transformer: A Heterogeneous 3D (H3D) Computing Platform for Transformer Model Acceleration on Edge Devices. *ACM Trans. Des. Autom. Electron. Syst.* 29, 1–19. (doi:10.1145/3649219)
14. Yang X, Yan B, Li H, Chen Y. 2020 ReTransformer: ReRAM-based processing-in-memory architecture for transformer acceleration. In *Proceedings of the 39th International Conference on Computer-Aided Design*, pp. 1–9. (doi:10.1145/3400302.3415640)
15. Zhou M, Xu W, Kang J, Rosing T. 2022 TransPIM: A Memory-based Acceleration via Software-Hardware Co-Design for Transformer. In 2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA), Seoul, Korea, pp. 1071–1085. (doi:10.1109/HPCA53966.2022.00082)
16. Mattson P, Cheng C, Coleman C, Damos G, Micikevicius P, Patterson D, Tang H *et al.* 2019 MLPerf Training Benchmark. *arXiv Preprint arXiv*
17. Liu Z, Li G, Cheng J. 2021 Hardware Acceleration of Fully Quantized BERT for Efficient Natural Language Processing. In 2021 Design, Automation & Test in Europe Conference & Exhibition (DATE), Grenoble, France, pp. 513–516. (doi:10.23919/DATE51398.2021.9474043)
18. Amirshahi A, Klein JAH, Ansaloni G, Atienza D. 2023 Tic-sat: tightly-coupled systolic accelerator for transformers. In *Proceedings of the 28th Asia and South Pacific Design Automation Conference*, pp. 657–663. (doi:10.1145/3566097.3567867)
19. Glint T, Jha CK, Awasthi M, Mekie J. 2023 Analysis of Conventional, Near-Memory, and In-Memory DNN Accelerators. In 2023 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Raleigh, NC, USA, pp. 349–351. (doi:10.1109/ISPASS57527.2023.00049)
20. Qin Y, Wang Y, Deng D, Zhao Z, Yang X, Liu L, Wei S, Hu Y, Yin S. 2023 FACT: FFN-Attention Co-optimized Transformer Architecture with Eager Correlation Prediction. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*,



- Orlando FL USA, pp. 1–14. New York, NY, USA. (doi:10.1145/3579371.3589057). <https://dl.acm.org/doi/proceedings/10.1145/3579371>.
21. Kim KM, Jeong YS, Bang IC. 2019 Thermal analysis of lithium ion battery-equipped smartphone explosions. *Eng. Sci. Technol. Int. J.* **22**, 610–617. (doi:10.1016/j.jestch.2018.12.008)
  22. Glint T, Prasad K, Dagli J, Gandhi K, Gupta A, Patel V, Shah N, Mekie J. 2023 Hardware-Software Codesign of DNN Accelerators Using Approximate Posit Multipliers. In ASPDAC, Tokyo Japan, pp. 469–474. New York, NY, USA. (doi:10.1145/3566097.3567866). <https://dl.acm.org/doi/proceedings/10.1145/3566097>.
  23. Glint T, Awasthi M, Mekie J. 2024 Hardware-Software Co-Design of a Collaborative DNN Accelerator for 3D Stacked Memories with Multi-Channel Data. In *29th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 454–459. (doi:10.1109/ASP-DAC58780.2024.10473935)
  24. Hynix SK. SK hynix partners with TSMC to strengthen HBM technological leadership. See <https://news.skhynix.com/sk-hynix-partners-with-tsmc-to-strengthen-hbm-technological-leadership/> (accessed 16 September 2024).
  25. Gohil V, Walia S, Mekie J, Awasthi M. 2021 Fixed-Posit: A Floating-Point Representation for Error-Resilient Applications. *IEEE Trans. Circuits Syst. II* **68**, 3341–3345. (doi:10.1109/tcsii.2021.3072217)
  26. Muralimanohar N, Balasubramonian R, Jouppi NP. 2009 CACTI 6.0: A Tool to Model Large Caches. *HP Labs.* **27**, 28. (doi:10.1109/MM.2008.2)
  27. Wu YN, Emer JS, Sze V. 2019 Accelergy: An Architecture-Level Energy Estimation Methodology for Accelerator Designs. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, Westminster, CO, USA, pp. 1–8. IEEE. (doi:10.1109/ICCAD45719.2019.8942149). <https://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=8931666>.
  28. Glint T. 2024 AxLaM: Energy-Efficient Accelerator Design for Language Models for Edge Computing. Zenodo. (doi:10.5281/zenodo.14196450)